

DTIC FILE COPY

2

AD-A219 276

**A PARALLEL DISTRIBUTED PROCESSING
APPROACH TO BEHAVIOR AND BIOLOGY
IN SCHIZOPHRENIA**

Technical Report AIP - 100

J. D. Cohen and D. Servan-Schreiber

Carnegie Mellon University,
University of Pittsburgh &
Stanford University

**The Artificial Intelligence
and Psychology Project**

Departments of
Computer Science and Psychology
Carnegie Mellon University

Learning Research and Development Center
University of Pittsburgh

DTIC
ELECTE
MAR 14 1990
S B D

Approved for public release; distribution unlimited.

90 03 12 076

**A PARALLEL DISTRIBUTED PROCESSING
APPROACH TO BEHAVIOR AND BIOLOGY
IN SCHIZOPHRENIA**

Technical Report AIP - 100

J. D. Cohen and D. Servan-Schreiber

Carnegie Mellon University,
University of Pittsburgh &
Stanford University

This research was supported by the Computer Sciences Division, Office of Naval Research and DARPA under Contract Number N00014-86-K-0678. Reproduction in whole or in part is permitted for purposes of the United States Government. Approved for public release; distribution unlimited.

DTIC
ELECTE
MAR 14 1990
S B D

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE

REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION Unclassified			1b. RESTRICTIVE MARKINGS	
2a. SECURITY CLASSIFICATION AUTHORITY			3. DISTRIBUTION / AVAILABILITY OF REPORT Approved for public release; Distribution unlimited	
2b. DECLASSIFICATION / DOWNGRADING SCHEDULE				
4. PERFORMING ORGANIZATION REPORT NUMBER(S)			5. MONITORING ORGANIZATION REPORT NUMBER(S)	
6a. NAME OF PERFORMING ORGANIZATION Carnegie-Mellon University		6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION Computer Sciences Division Office of Naval Research	
8a. ADDRESS (City, State, and ZIP Code) Department of Psychology Pittsburgh, Pennsylvania 15213			7b. ADDRESS (City, State, and ZIP Code) 800 N. Quincy Street Arlington, Virginia 22217-5000	
8a. NAME OF FUNDING / SPONSORING ORGANIZATION Same as Monitoring Organization		8b. OFFICE SYMBOL (If applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER N00014-86-K-0678	
8c. ADDRESS (City, State, and ZIP Code)			10. SOURCE OF FUNDING NUMBERS p400014201/7-4-86	
			PROGRAM ELEMENT NO N/A	PROJECT NO N/A
			TASK NO. N/A	WORK UNIT N/A
11. TITLE (Include Security Classification) "A Parallel Distributed Processing Approach to Behavior and Biology in Schizophrenia"				
12. PERSONAL AUTHOR(S) Jonathan Cohen and David Servan-Schreiber				
13a. TYPE OF REPORT Technical	13b. TIME COVERED FROM 86Sept15 to 91Sept1	14. DATE OF REPORT (Year, Month, Day) 1989-October-01	15. PAGE COUNT	
16. SUPPLEMENTARY NOTATION				
17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)	
FIELD	GROUP	SUB-GROUP	Schizophrenia; Connectionist; Attention; Frontal Lobes; Dopamine; Cognitive Neuroscience; Language	
19. ABSTRACT (Continue on reverse if necessary and identify by block number) SEE REVERSE SIDE				
20. DISTRIBUTION / AVAILABILITY OF ABSTRACT <input type="checkbox"/> UNCLASSIFIED/UNLIMITED <input checked="" type="checkbox"/> SAME AS RPT <input type="checkbox"/> DTIC USERS			21. ABSTRACT SECURITY CLASSIFICATION	
22a. NAME OF RESPONSIBLE INDIVIDUAL Dr. Alan L. Mayrovitz			22b. TELEPHONE (Include Area Code) (202) 696-4302	22c. OFFICE SYMBOL N00014

ABSTRACT

Schizophrenia is marked by a variety of cognitive and biological abnormalities. In the first part of this paper we describe schizophrenic cognitive deficits in three experimental tasks which tap attention and language processing abilities. We also review biological disturbances that have been reported involving the frontal lobes and the mesocortical dopamine system. In the second part of the paper we present three computer models, each of which simulates normal performance in one of the cognitive tasks described initially. These models were developed within the connectionist (or parallel distributed processing) framework. At the behavioral level, the models suggest that a disturbance in the processing of context can account for schizophrenic patterns of performance in both the attention and language-related tasks. At the same time, the models incorporate features of biological computation that address the biological processes underlying cognitive deficits. All three models incorporate a mechanism for processing context that can be identified with frontal lobe function, and a parameter that corresponds to the effects of dopamine on frontal cortex. A disturbance in this parameter is sufficient to account for schizophrenic patterns of performance in all three of the cognitive tasks simulated. Thus, the models offer an explanatory mechanism linking performance deficits to a disturbance in the processing of context which, in turn, is attributed to a reduction of dopaminergic activity in prefrontal cortex. In the General Discussion, we consider the implications of these models for our understanding of both normal and schizophrenic cognition. We conclude with a discussion of some of the general issues surrounding the modelling endeavor itself.



Accession For	
NTIS GRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	

Table of Contents

I. Introduction	1
II. Cognitive and Biological Deficits in Schizophrenia	2
A. Cognitive Deficits	2
1 Schizophrenic Deficits of Attention	2
2 Schizophrenic Language Deficits	6
B. Biological Deficits	9
1 Frontal Cortex and Schizophrenia	10
2 Dopamine and Schizophrenia	13
C. Summary	14
III. Connectionist Simulations of Biological and Behavioral Disturbances	15
A. The Connectionist Framework	15
B. Simulation of the Physiological effects of Dopamine	17
C. A connectionist Model of Selective Attention (the Stroop Effect)	19
D. Simulation of the Continuous Performance Test (CPT)	22
E. Simulation of Context-Dependent Lexical Disambiguation	24
IV. General Discussion	29
1. Attention and Context	29
2. Disturbance in the Processing of Context	30
3. Generalized versus Specific Deficits	32
4. Frontal Cortex, Dopamine and the Processing of Context	32
5. Biological Disturbances in Schizophrenia	35
6. Comparison with Other Models of Schizophrenia	36
7. The Role of Computational Modelling	38
V. Conclusion	39
Authors' Notes	40
References	41

I. Introduction

Schizophrenia is marked by a wide variety of behavioral deficits, including disturbances of attention, language processing and problem solving. At the same time, there is accumulating evidence regarding biological abnormalities in schizophrenia, including disturbances in specific neurotransmitter systems (e.g., dopamine) and anatomic structures (e.g., frontal cortex). Unfortunately, however, there is still only a poor understanding of how the various behavioral deficits relate to one another, or how these deficits arise from disturbances at the biological level. In the absence of a theory of information processing linking behavior to neural events, it is difficult to propose explanations integrating anatomical, physiological and behavioral observations.

In an effort to address this problem, we have drawn upon the recent development of the connectionist framework. This framework provides a means for building computer simulation models of cognitive phenomena. However, connectionist models are distinct from other computer modelling efforts in cognitive science in their use of information processing mechanisms that incorporate important features of biological computation. Using this framework, we have begun to develop models that explore the effect of biologically relevant variables on behavior in schizophrenia. In this paper, we present three such models.

At the behavioral level, we focus on schizophrenic disturbances of attention and language. We describe a set of connectionist models that simulate both normal and schizophrenic patterns of performance in three experimental tasks that tap attentional and language processing abilities. The models make use of a common set of information processing mechanisms, and show how a number of seemingly disparate observations about schizophrenic behavior can all be related to a single functional deficit: a disturbance in memory for context.

The models also suggest a direct link between this functional deficit and specific biological abnormalities in schizophrenia. In particular, our models address disturbances in two systems that have consistently been implicated in the pathophysiology of schizophrenia: the prefrontal cortex and the mesocortical dopamine system. We will argue that one component of our models implements the function of the prefrontal cortex: maintenance of information necessary for the selection of action (i.e., memory for context). We will also argue that changes in a particular parameter of the models (the gain parameter) corresponds to the effects of dopamine on cortical neurons: a modulation of their sensitivity to afferent input. The models show how a disturbance of this parameter localized to the part of the network corresponding to prefrontal cortex can explain schizophrenic patterns of behavior in three separate experimental tasks. Taken together, the models suggest that disturbances of attention and language in schizophrenia can be accounted for in terms of a decrease in dopaminergic effects in frontal cortex, and provide a precise account of the way in which the behavioral effects arise from biological disturbances.

A primary goal of our efforts is the integration of biological and behavioral findings in schizophrenia research. At the same time, we recognize that it is not possible to address all of the phenomena relevant to this complex illness. In this paper we focus on a subset of the pathological phenomena associated with schizophrenia. In so doing, we set aside a number of important behavioral phenomena that are characteristic of this disease (e.g., hallucinations,

delusions, and disturbances of affect), as well several biological abnormalities (e.g., cellular disorganization of the hippocampus, and disturbances of non-dopaminergic neurotransmitter systems). The models we present are not meant to be a complete theory of schizophrenia. Rather, they are intended to provide insight into the mechanisms that underlie certain aspects of this illness. Nevertheless, we believe that a more precise account of attentional and language disturbances will be valuable for tackling more complex phenomena, such as hallucinations and delusions; and that, at the biological level, our simulations of prefrontal cortex and the effects of dopamine will provide a starting point for simulating interactions between these systems and others that have been implicated in schizophrenia. More generally, we hope that our attempt to bridge the traditional gap between biological and behavioral research will provide a useful example for similar efforts in other areas of research.

We begin by reviewing the data concerning cognitive and biological deficits in schizophrenia that are relevant to our hypotheses. These are diverse bodies of literature. To provide coherence to our review, we will point to interpretations of the data that indicate the relationship between the different empirical phenomena that we are interested in. These interpretations arose from our work with a set of connectionist models that can be used to simulate the phenomena. In the second part of the paper we describe these models. We then consider the implications of these models — and the interpretations of the empirical data they provide — in the General Discussion.

II. Cognitive and Biological Deficits in Schizophrenia.

A. Cognitive Deficits

A large number of experiments have revealed schizophrenics deficits in attention and language processing tasks. Below, we consider schizophrenic performance in three tasks that are representative of these domains of processing, and in which schizophrenics show characteristic deficits. In particular, we focus on the role that memory for context plays in these tasks. By “memory for context”, we mean memory for information that is necessary to select an appropriate response, but that is not actually part of the *content* of the response. Furthermore, we want to distinguish this type of memory (which can be thought of as a type of short term, or working memory) from long term memory (such as that involved in associative, or reinforcement learning). The findings we review suggest that a degradation in representing and maintaining context underlies many of the deficits in attention and language processing observed in schizophrenics.

1 Schizophrenic Deficits of Attention

A fundamental aspect of human attention is the ability to act on one set of stimuli, even when other, possibly more compelling stimuli are available. This is exhibited, for example, in our ability to pick out a single instrument in an orchestral arrangement, to identify a face in a crowd, and to concentrate on a difficult mental problem on the bus ride to work, screening out stimuli from the environment. Investigators who have focused on the phenomenology of schizophrenia have often reported that patients appear unable to screen out irrelevant stimuli from the environment, indicating a deficit in selective attention (e.g., McGhie & Chapman,

1961; McGhie, 1970; Lang & Buss, 1965; Garnezy, 1977). Perhaps the experimental task that has most commonly been used to study selective attention — in normal subjects — is Stroop task (Stroop, 1935; for reviews, see Dyer, 1973 and MacLeod, 1989). There have also been several applications of this task to the study of schizophrenics.

a. The Stroop Task

The Stroop task consists of two subtasks. In one, subjects name the color of the ink in which a word is printed. In the other, subjects read the word aloud while ignoring ink color. Three types of stimuli are used: conflict stimuli, in which the word and the ink color are different (e.g., the word RED in green ink);¹ congruent stimuli, in which they are the same (e.g., the word RED in red ink); and control stimuli. The control stimuli for word reading are typically color words printed in black ink; for the color naming they are usually a row of XXXX's printed in a color. The subjective experience of performing this task is that word reading is much easier, and there is no difficulty in ignoring the color of the ink. In contrast, it is much harder to ignore the word when the task is to name ink color.




Figure 1. Performance in the standard Stroop task (after Dunbar & MacLeod, 1984). Data are average reaction times to stimuli in each of the three conditions of the two tasks.

These phenomena are reflected in the time it takes for subjects to respond to stimuli of each type (see Figure 1). Three basic effects are commonly observed: 1) word reading is faster than color naming; 2) ink color has no effect on the speed of word reading; and 3) words have a large effect on color naming (slowing it when the word conflicts with the color to be named, and speeding it when the word agrees). For example, subjects are slower to respond to the color red when the word GREEN is written in red ink, than when the word RED or a series of

¹ Throughout this discussion, references to word stimuli will appear in upper case (e.g., RED), references to color stimuli will appear in lower case (red), and references to potential responses will appear in quotation marks ("red").

X's appear in red ink. Thus, subjects are less able to selectively attend to colors (i.e., ignore words) than the reverse. If schizophrenics suffer from a deficit in selective attention, then they should show a larger Stroop effect; that is, they should be less able to ignore word information, and show a *greater interference effect*.

Table 1.
Performance of Normal and Schizophrenic Subjects
in Two Studies Using the Stroop Task.*

		Wapner & Krus (1960)	Abramczyk, Jordan & Hegel (1983)
Normal Controls	Word reading	39	43
	Color naming	57	60
	Color naming Interference	98	100
Schizophrenics	Word reading	57	50
	Color naming	78	77
	Color naming Interference	151	140

* Both studies used the original form of the Stroop task, in which subjects are given three cards, one with color words written in black ink (word reading), one with color patches or XXX's print in different colors (color naming), and one with color words each written in a conflicting ink color (Color naming interference). Data are the average number of seconds subjects took to respond to all of the stimuli on each type of card.

Table 1 reports data from two empirical studies comparing normal and schizophrenic performance in the Stroop task.² Performance of normal control subjects conformed with the standard findings in this task: subjects were faster at reading words than naming colors, and words interfered with color naming. Schizophrenics also showed this pattern of results. However, in both studies schizophrenics differed significantly from controls in two important ways: 1) schizophrenics showed an *overall* slowing of responses; and 2) they showed a

² To our knowledge, there are only 4 studies reported in the literature that tested schizophrenics using the standard Stroop task (Wapner and Krus, 1960; Grand et al., 1975; Abramczyk et al., 1983; Mirsky et al., 1983). Only three of these report reaction times, and one involved only four subjects (Mirsky et al., 1983). The data for these four subjects, while statistically unreliable, conformed to the overall pattern of our predictions. That is, subjects showed disproportionate amounts of interference. Interestingly, this worsened when they were taken off of medication. The data for the two remaining studies are presented in Table 2.

statistically *disproportionate* slowing of responses in the interference condition of the color naming task. On first consideration, the latter finding would appear to be predicted by a schizophrenic deficit in selective attention. However, this interpretation has been called into question: the general non-specific slowing of performance may be responsible for the additional interference, rather than a specific attentional deficit (see Chapman & Chapman, 1978 for a discussion of the general issue of differential vs. generalized deficit). Because of the general nature of these competing claims, it is not possible to decide which, if either of these accounts is correct. One purpose of the simulations we present below is to commit each of these hypotheses to specific information processing mechanisms, and compare their ability to account for the data.

b. The Continuous Performance Test

While the Stroop task stands at the center of research on attention in normal populations, other tasks have been used more extensively to study attentional deficits in schizophrenics. One of these is the continuous performance test (CPT — Rosvold, Mirsky, Sarason, Bransome & Beck, 1956). In this task, subjects are asked to detect a target event among a sequence of briefly presented stimuli, and to avoid responding to distractor stimuli. The target event may be the appearance of a single stimulus (e.g., detect the letter 'X' appearing in a stream of other letters), or a stimulus appearing in a particular context (e.g., respond to 'X' only when it follows 'A', or respond to the consecutive repetition of any letter). The percentage of correctly reported targets (hits) and of erroneous responses to distractors (false alarms) give a measure of the subjects' signal detection ability; that is, their ability to discriminate between target and non-target events, independent of their response bias (cf. Green & Swets, 1966, for a description of signal detection theory). Schizophrenic patients typically show lower hit rates and similar or higher false alarm rates compared to normal subjects and patient controls (e.g., Kornetsky, 1972; Spohn, Lacoursiere, Thomson & Coyne, 1977; Nuechterlein 1984), indicating poorer signal detection ability. This is especially true when the task makes high processing demands (e.g., when stimuli are degraded or when memory of the previous stimulus is necessary). The fact that schizophrenics show impaired signal detection performance, independent of response bias, indicates that their poorer performance is not due simply to lack of motivation (e.g., ignoring the task altogether) or to arbitrary responding (Swets & Sewell, 1963). Rather, this pattern of results indicates that schizophrenics have difficulty in discriminating between stimuli and distractors. This impairment could result from an inability to make effective use of context necessary to perform the task. This interpretation is supported by schizophrenics' performance in versions of the task in which the response to the current stimulus is contingent on previous stimuli. For example, in the 'CPT-double' a target event consists of two consecutive identical letters. Memory for the previous letter provides the necessary context to evaluate the significance of the current letter. Schizophrenics perform especially poorly in this and similar versions of the task (Nuechterlein, 1984).

Impaired CPT performance is a highly reliable and stable marker of schizophrenia. It is found across a large spectrum of the clinical and subclinical presentations of the disease. For example, several studies have shown that performance remains impaired even when clinical status changes from the acute hospitalization phase to clinical remission, and when patients are administered neuroleptic medications (Asarnow & MacCrimmon, 1978; Spohn et al., 1977). Siblings and offspring of schizophrenic patients tend to display a similar deficit on the CPT compared to relatives of patients suffering from other psychiatric disorders (Rutschmann, Cornblatt & Erlenmeyer-Kimling, 1977; Erlenmeyer-Kimling & Cornblatt, 1978; Nuechterlein, 1983). This finding rules out an explanation of impaired CPT performance based on 'nuisance variables' such as repeated hospitalizations, or side effects of past neuroleptic treatments.

Finally, CPT performance of schizophrenics seems to be related to the clinical activity of the disease. Performance improves with long-term neuroleptic therapy (Spohn et al., 1977), and particularly so in patients who also experience greater clinical benefits from treatment (Kornetsky, 1972).

Other measures of attention. In addition to the Stroop task and the CPT, there are a number of other information processing paradigms in which schizophrenics exhibit performance deficits, including the span of apprehension task (Neale, 1971), studies of dichotic listening (Spring, in press; Wielgus & Harvey, 1988), and a variety of reaction time tasks (see Nuechterlein, 1977 for a review of the early literature, and Borst & R. Cohen, 1989 and R. Cohen, Borst & Cohen, 1989 for more recent work). The prevailing interpretation of these data is that they reflect a disturbance of attention in schizophrenia. For example, Shakow's (1962) original formulation in terms major and minor sets is still frequently referred to: normal subjects are able to adopt a "major set" that takes account of all the various factors involved in performing the task; schizophrenics are unable to do so, relying instead on a "minor set" that takes account of only a limited set of factors (e.g., the most recent events). Shakow argued that these findings are indicative of "the various difficulties created by *context* [*sic*]... It is as if, in the scanning process which takes place before the response to a stimulus is made, the schizophrenic is unable to select out the material relevant for optimal response." (Shakow, 1962). As yet, however, there is no generally accepted understanding of the specific information processing mechanisms that are involved in maintaining an attentional set, and that explain its relationship to the processing of context in schizophrenia.

2 Schizophrenic Language Deficits

Schizophrenics also show poor use of context in language processing. Chapman, Chapman & Miller (1964) first described this in their study of schizophrenics' interpretation of lexical ambiguities. They found that schizophrenics tended to interpret the strong (dominant) meaning of a homonym used in a sentence, even when context provided by the sentence mediated the weaker (subordinate) meaning. For example, given the sentence "The farmer needed a new *pen* for his cattle," schizophrenics interpreted the word "pen" to mean writing implement more frequently than control subjects. They did not differ from control subjects in the number of unrelated meaning responses that were made (e.g., interpreting "pen" to mean "fire truck"), nor did they differ in the number of types of errors that they made when the strong meaning of the homonym was correct. These findings have been replicated in a number of studies (e.g., Benjamin & Watt, 1969; Blaney, 1974; Strauss, 1975; Cohen, Targ, Kristoffersen & Spiegel, 1989).

Other studies of language performance also support the view that schizophrenics make poor use of context, including those using cloze analysis (guessing the words deleted from a transcript of speech — e.g., Salzinger, Portnoy & Feldman, 1964; Salzinger, Portnoy, Pisoni & Feldman, 1970), speech reconstruction (ordering sentences which have been randomly rearranged — Rutter, 1979), and cohesion analysis (examining the types of references used in speech — e.g., Rochester & Martin, 1979; Harvey, 1983). (For reviews of this literature see Maher, 1972; Schwartz, 1982; Cozolino, 1983; and Cohen et al, 1989)

As in attentional tasks, it appears that schizophrenics suffer from a restriction in the sequential range over which contextual interactions occur. Thus, for example, Salzinger et al. (1964, 1970) found that schizophrenics and normals performed comparably well in "clozing" speech (i.e., guessing the word which was deleted from a sample of normal speech) when contextual

cues were local (e.g., when the missing word was surrounded by only two or three words). However, when a missing word was surrounded by larger numbers of words, normals improved in their ability to predict the word, whereas schizophrenics did not. This suggested that normals were able to make use of the additional context provided by cues more distal to the word, while schizophrenics could not. Conversely, Salzinger also showed that it is easier for normals to cloze small segments of schizophrenic speech than larger ones. This implies that broader segments of schizophrenic discourse do not add contextual constraint, presumably because schizophrenics produce contextual references which span more limited segments of speech. Based on these data, Salzinger proposed an immediacy hypothesis which stated that "the behavior of schizophrenic patients is more often controlled by stimuli which are immediate in their spatial and temporal environment than is that of normals" (Salzinger, 1971 — p. 608).

We recently tested the idea that schizophrenics are restricted in the temporal range over which they process context in language (Cohen et al., 1989). We designed a task, similar to the one used by Chapman and his colleagues (1964), in which subjects interpreted lexical ambiguities used in sentences. However, in our study we manipulated temporal parameters of the task. Subjects were presented with sentences made up of two clauses; each clause appeared one at a time on a computer screen. One clause contained an ambiguous word in neutral context (e.g., "you need a PEN"), while the other clause provided disambiguating context (e.g., "in order to keep chickens" or "in order to sign a check"). Clauses were designed so that they could be presented in either order: context first or context last. The ambiguity in each sentence always appeared in capital letters, so that it could be identified by the subject. Ambiguities were used which were known to have a strong (dominant) and a weak (subordinate) meaning,³ and a context clause was designed for each of the two meanings.

Subjects were presented with 54 sentences (one for each ambiguity) distributed across three conditions: a) weak meaning, context last; b) weak meaning, context first; c) strong meaning, context first. For example, the ambiguity "pen" would have appeared in one of the three following conditions:

(A) without a PEN [clear screen / pause] you can't keep chickens

— or —

(B) you can't keep chickens [clear screen / pause] without a PEN

— or —

(C) you can't sign a check [clear screen / pause] without a PEN

[clear screen / pause]

The meaning of the word in capital letters is:

a writing implement	(dominant meaning)
a fenced enclosure	(subordinate meaning)
a kind of truck	(unrelated meaning)

³ These were normed in a population of undergraduate students.

Following presentation of the two clauses comprising a sentence, subjects were presented with a list of meanings; they were asked to pick the meaning that best corresponded to the meaning of the ambiguity as it was used in the sentence (see example above). In each case, two of the response choices were related to the ambiguity, while a third was unrelated to either of its meanings.

The results of this study (shown in Figure 2) corroborated both the Chapmans' original findings, and the explanation of their findings in terms of a restriction in the temporal range over which schizophrenics are able to use context. Thus, schizophrenics made significantly more dominant meaning errors than did controls when the weak meaning was correct. However, this only occurred when the context came *first* (condition *B* above). When context came last, schizophrenics did not differ from patient controls. Nor, as the Chapmans found, did the groups differ in the number of weak meaning choices made when the strong meaning was correct, or in the number of unrelated meaning choices made in any condition. Thus, schizophrenics appear to have had difficulty using context, but only when it was temporally remote (i.e., came first), and not when it was more recently available (i.e., came last). This effect is consistent with Salzinger's immediacy hypothesis. Moreover, it suggests that the impairment observed in language tasks may be of a similar nature to the impairments observed in attentional tasks: difficulty in remembering and using context to control action. This impairment is evident when the contextually appropriate behavior is — in the absence of context — subordinate to a more dominant response tendency, as in the case of the ambiguities used in the Chapmans' and our study.

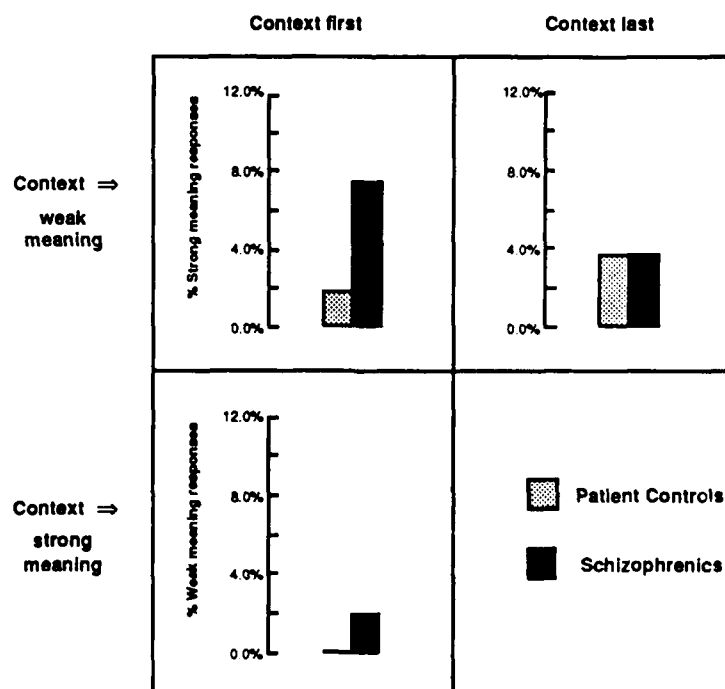


Figure 2. Medians for the rates of incorrect meaning errors for schizophrenics and patient controls. Due to the low overall rate of unrelated errors, and no significant differences between groups in this type of error, these data are not shown.

B. Biological Deficits

In parallel to research on schizophrenic information processing deficits, there has also been an intensive investigation of the biological abnormalities in schizophrenia. A number of different anatomic and physiological systems have been implicated, however little research has addressed how disturbances of these systems can lead to the types of information processing deficits that are observed. In this section we review data concerning the role of prefrontal cortex and the mesocortical dopamine system in information processing, and abnormalities of these systems in schizophrenia. In the next section we will describe a set of information processing mechanisms that can simulate important functions performed by these systems, and we will show how a specific disturbance in these mechanisms can lead to the performance deficits observed in schizophrenics.

1 Frontal Cortex and Schizophrenia

a. Function of Prefrontal Cortex

It is commonly accepted that the frontal lobes are involved in the planning and sequencing of complex actions (e.g., Luria, 1966; Shallice, 1982). Memory for context is obviously an important determinant of such goal-oriented behavior. The actions associated with a particular goal may, in other contexts, be relatively infrequent or "weak" behaviors. Such actions require the maintenance of an internal representation of the goal — or of knowledge related to it — to favor their execution, and to suppress competing, possibly more compelling behaviors. For example, we have all struggled with the urge to scratch a mosquito bite. Resisting this urge relies on actively accessing the knowledge that if the bite is left alone it will resolve more quickly. This knowledge can be thought of as the context needed to control behavior, and it must be actively maintained or the prepotent response tendency (scratching the bite) will prevail (e.g., as they do during sleep or while absorbed in another activity).

Recent studies have begun to supply direct evidence that frontal areas are involved in maintaining context for the control of action. Fuster (1980), Goldman-Rakic (1985) and Diamond (1989a; 1989b; Diamond & Doar, 1989) have all reported experimental data which show that prefrontal cortex is needed to perform tasks involving delayed responses. These studies have provided insights at both the biological and behavioral levels. For example, using single unit recording techniques, Fuster (1980) and Goldman-Rakic's group (Brozoski, Brown, Rosvold & Goldman, 1979) have observed cells in prefrontal cortex that are specific to a particular stimulus and response, and that remain active during the delay between presentation of the stimulus and execution of the response. They have argued that neural patterns of activity are maintained in prefrontal cortex which encode the temporary information needed to guide a response. Diamond (e.g., 1985; 1989c) has emphasized that prefrontal memory is required, in particular, to overcome competing, prepotent response tendencies in order to mediate a contextually relevant — but otherwise weaker — response. She cites extensive data from lesion studies in adult monkeys and from developmental studies in human and monkey infants that use a variety of behavioral tasks (including object retrieval, visual paired comparisons, delayed response, and the A-B task). Results from these studies suggest that prefrontal cortex is directly involved in maintaining representations that inhibit reflexive or habitually reinforced

behaviors to attain a goal. This is most clearly demonstrated in the A \overline{B} task (pronounced "A not B").

In the A \overline{B} task (Piaget, 1954[1937]), subjects observe a desired object being hidden at one of two locations which are identical in appearance. Attention is then drawn away from both locations for a specified delay, after which they are allowed to retrieve the object. On subsequent trials, the object is hidden at the same location until it is successfully retrieved some number of times in a row. Then it is hidden at the other location. Thus, there are two variables of interest in this task: the duration of the delay imposed between hiding and retrieval, and the location at which the object is hidden (same or different from the previous trial). Normal adult monkeys and five year old human children can successfully retrieve the object — independent of location — with delays between hiding and retrieval of two minutes or more. Monkeys with lesions of prefrontal cortex, as well as human infants younger than 6 months can perform the task successfully only if there is *no* delay. At delays of two seconds or more, their performance degrades. However, it does so in a systematic way: when the location remains the same, they perform acceptably; but when it is switched subjects return to the location at which the object was last hidden, even though they see it being hidden at the new location. This pattern of errors is specific to human infants, and to monkeys with lesions of prefrontal cortex, and is not found with lesions of the hippocampus or parietal lobes. In the latter case, performance is at chance (i.e., is independent of location), while with hippocampal lesions performance is normal up to delays of 30 seconds, after which it drops to chance. The interpretation of these findings is that subjects lacking prefrontal cortex are unable to hold in memory a context representation (the location of the hidden object) required to inhibit a dominant response (return to the most recently rewarded location). Note that this interpretation draws a distinction between memory for context — preceding events that do not themselves trigger action — and memory for previously rewarded stimulus-action pairs (i.e., reinforcement). These two forms of memory are assumed to be supported by different neural structures, with prefrontal cortex involved only in the former. It is precisely because lesions of prefrontal cortex affect one memory system (memory for context) and not the other (reinforcement) that a perseverative pattern of performance can occur. As we have noted, lesions which involve other areas (e.g., hippocampus) result in random rather than perseverative behavior (e.g., Diamond, 1989b).

The performance deficits observed for infants and frontally lesioned monkeys on delay tasks is very similar to those observed for human frontal lobe patients on the Wisconsin Card Sort Task (WCST — Grant & Berger, 1948). In this task, subjects are presented with a series of cards containing figures that vary in three ways: shape, color and number. They are asked to sort the cards into piles according to a rule that the experimenter has in mind (e.g., separate the cards by color). However, subjects are not explicitly told the rule for sorting; rather, for each card they are given feedback as to whether or not they have sorted the card properly. Normal subjects discover the rule quickly. Once they have demonstrated that they know it (i.e., by making a minimum number of consecutive correct sorts) the experimenter switches the rule, and the subject is required to discover the new rule. Patients with damage to the frontal lobes do poorly on this task (e.g., Milner, 1963; Nelson, 1976; Robinson, Heaton, Lehman & Stilson, 1980). While they are able to discover the first rule without too much difficulty, they have a hard time switching to a new rule: they continue to sort according to the old one. Strikingly, some subjects have shown this perseveratory behavior despite their ability to verbalize the correct new rule, or to comment on the fact that they know what they are doing is wrong. This behavior is qualitatively similar to that observed for infants and frontally lesioned monkeys on delay tasks. In both sets of tasks, subjects are required to overcome the tendency to repeat response patterns that were correct on previous trials. Note that in both sets of tasks,

subjects with poor prefrontal function are not impaired in their ability to learn the basic elements of the task. Rather, they are impaired in their ability to use current context to *override* the effects of prior experience in the task. This characterization of frontal lobe function jibes with the clinical characterization of frontal lobe deficits as a "disinhibition syndrome" (Stuss & Benson, 1984). It is also consistent with the difficulties that have been observed for frontal lobe patients in performing the Stroop task (Perret, 1974) and similar tasks in clinical use (e.g., the "go-no-go" paradigm) that require the inhibition of a dominant response tendency. Finally, physiological measures have begun to provide converging evidence for this hypothesis.

In studies of regional cerebral blood flow (rCBF), Weinberger and his collaborators (Weinberger, Berman & Zec, 1986; Berman, Illowsky & Weinberger, 1988; Weinberger, Berman & Chase, 1988) have shown that in normal subjects metabolism increases selectively in prefrontal cortex during performance of the WCST. Moreover, WCST performance was correlated with the increase in prefrontal cortex metabolism relative to other cortical areas. This finding corroborates the result of neuropsychological studies that link WCST performance with frontal lobe function (e.g., Nelson, 1976; Robinson et al., 1980). Weinberger's group also showed that not all cognitive tasks requiring effort and concentration are accompanied by such an increase in prefrontal activity. For example, during the Raven Progressive Matrices test — in which the task-relevant information is visually available at all times — metabolism increased in parietal and occipital areas but not in frontal areas.

If frontal cortex is generally involved in maintaining context for the selection of action and, as we have argued, CPT performance relies on this ability, then changes in frontal metabolism should be observed during this task as well. Indeed, there is data to support this supposition. R. M. Cohen and his colleagues (Cohen et al., 1987; Cohen, Semple, Gross, Holcomb, Dowling & Nordahl, 1988) studied the involvement of PFC during an auditory-discrimination version of the CPT, using positron emission tomography (PET) to measure regional metabolism. They found an increase in prefrontal metabolism in normal subjects with this task as well, and a correlation between CPT performance and prefrontal metabolism: subjects who made more commission errors (false alarms) showed less of an increase in prefrontal metabolism. We should note that not all studies examining frontal lobe function during CPT performance have yielded positive results (e.g., Berman, Zec & Weinberger, 1986). However, differing results may be attributable to differences in the actual tasks and conditions that were run. We will return to this issue below, in the General Discussion.

In summary, evidence from attentional tasks (e.g., CPT and Stroop), problem solving tasks (e.g., WCST and A B), and from physiological studies suggest that areas of the frontal cortex support memory for information needed for response selection. Failure of this memory system is most apparent when experimental tasks involve competing, prepotent responses. These may have developed during the task itself (as in the WCST and A B), or they may have existed prior to the experiment (e.g., the Stroop task). Failure in the prefrontal memory system manifests as a bias toward prepotent, but task-inappropriate response tendencies. The data reviewed earlier concerning schizophrenic performance deficits fits with this profile: an insensitivity to context — particularly when memory for context is involved — and a dominant response tendency. This suggests that disturbances of the frontal lobes may be involved in schizophrenia.

b. Frontal Deficits in Schizophrenia

The idea that frontal lobe function is impaired in schizophrenia is not new. Kraepelin, who first defined the illness dementia praecox that we now refer to as schizophrenia, wrote:

On various grounds it is easy to believe that the frontal cortex, which is especially well developed in man, stands in closer relation to his higher intellectual abilities, and that these are the faculties which in our patients invariably show profound loss... (Kraepelin, 1950 — p. 219).

Schizophrenics show typical frontal lobe deficits on standard neuropsychological tests, including the Wisconsin Card Sort task (e.g., Malmö, 1974; Kolb & Whishaw, 1983) and the Stroop task (e.g., Wapner & Krus, 1960; Abramczyk, Jordan & Hegel, 1983) (see Kolb & Whishaw, 1983 for a review). Recently, direct evidence has come from imaging and electrophysiological techniques. Ingvar and Franzen (1974; Franzen & Ingvar, 1975) and Gur et al. (1983; 1985) have reported abnormal perfusion of frontal areas in schizophrenics using positron emission tomography (PET), and Buchsbaum et al. (1982) have found abnormalities of glucose utilization localized to similar areas. Andreasen et al. (1986) have reported computerized tomographic (CT) images which indicate frontal lobe atrophy in schizophrenics, and other data suggest that ventricular enlargement in schizophrenics (e.g., Weinberger, Bigelow, Kleinman, Klein, Rosenblatt & Wyatt, 1980; Andreasen et al., 1986) is associated with frontal lobe atrophy (Morihsa & McAnulty, 1985). Farkas and colleagues (Farkas, Wolf, Jaeger, Brodie, Christman & Fowler, 1984) have demonstrated a correlation between abnormal structure (CT) and perfusion (PET) of the frontal lobes, while Morihsa and McAnulty (1985) showed a correlation between structural (CT) and electrophysiological abnormalities.

Recent studies have begun to examine the relationship between physiological and behavioral disturbances of frontal lobe function in schizophrenics. Weinberger et al. (1986) have demonstrated abnormal perfusion of the frontal lobes during performance of the WCST. Similarly, R. M. Cohen et al. (1987; Cohen, Semple, Gross, Nordahl, et al., 1988) have shown that schizophrenics fail to show the normal pattern of increased frontal lobe perfusion during performance of a variant of the CPT. Thus, the work emerging in this area suggests that anatomic and physiological deficits of frontal cortex are associated with the behavioral deficits that have been observed for schizophrenics.

2 Dopamine and Schizophrenia

The hypothesis that frontal lobe dysfunction is involved in schizophrenia fits well with the prevailing neurochemical and psychopharmacological data concerning this illness. The frontal cortex is a primary projection area for the mesocortical dopamine system, a disturbance of which has consistently been implicated in schizophrenia (e.g., Meltzer & Stahl, 1976; Nauta & Domesick, 1981). The dopamine hypothesis is one of the most enduring biological hypotheses concerning schizophrenia. Evidence for this hypothesis comes from a variety of sources. Perhaps the strongest argument is the chemical specificity of the neuroleptics, which are used to treat the symptoms of schizophrenia. *In vitro* studies have demonstrated that neuroleptics have a specific affinity for dopamine binding sites, and that this affinity is correlated with their clinical potency (Snyder, 1976; Creese, Burt & Snyder, 1976; B. Cohen,

1981). Furthermore, drugs that increase DA activity in the CNS — such as amphetamines and L-dopa — exacerbate symptoms in psychotic patients, and may induce psychosis in non-psychotic individuals (e.g., Snyder, 1972; Janowsky, Huey, Storms & Judd, 1977; Janowski & Rich, 1979). Studies of the plasma (Bowers, Heninger & Sternberg, 1980; Pickar et al., 1984) and cerebrospinal fluid (Sedvall, Fyro, Nyback, Wiesel & Wode-Helgodt, 1974) of schizophrenics have revealed elevated levels of dopamine metabolites. Finally, several post-mortem studies have found DA receptors in schizophrenic populations compared to matched controls (e.g., Cross, Crow & Owen, 1981), and this elevation is reliably correlated with the previous experience of hallucinations and delusions (Crow et al., 1984). Post-mortem studies have also revealed elevated levels of brain DA and DA metabolites in schizophrenics compared to matched controls (e.g., Bird, Barnes, Iversen, Spokes, Mackay & Shepherd, 1977).

While different investigators have argued that central dopamine activity is either reduced or increased in schizophrenia, one hypothesis is that both conditions may occur (either within or across individuals), and that each is associated with a different psychopathological profile. For example, Crow (1980) has suggested that the symptoms of schizophrenia can be divided into two subtypes, one of which reflects dopamine overactivity (positive symptoms — e.g., hallucinations and delusions) and the other that reflects dopamine underactivity (negative symptoms — e.g., avolition, amotivation and withdrawal). Several authors have argued that it is the negative symptoms of schizophrenia that are most often associated with hypofrontality (e.g., Levin, 1984; Andreasen et al., 1986). This is consistent with mounting evidence that mesocortical dopamine activity in frontal cortex is directly related to cognitive function, and that a reduction of this activity can produce many of cognitive deficits observed in schizophrenics. Thus, McCulloch and his colleagues (McCulloch, Savaki, McCulloch, Jehle & Sokoloff, 1982) have shown that activation of mesocortical dopaminergic neurons increases metabolic activity in the frontal cortex of animals. Conversely, lesions of the same dopamine projections reduce frontal metabolism and impair cognitive functions usually associated with frontal cortex, such as the execution of search strategies or delayed-alternation tasks (Oades, 1981; Simon, Scatton & Le Moal, 1980). For example, Rhesus monkeys could not perform a delayed-alternation task following selective destruction of DA terminals in prefrontal cortex (Brozoski et al., 1979). This deficit was as severe as that following full surgical ablation of the same area of cortex. Moreover, performance recovered almost entirely with DA agonists such as L-Dopa and apomorphine. Finally, studies of human patients suffering from Parkinson's disease — in which DA function is markedly impaired — provided similar evidence: even when these patients did not display clinically significant cognitive deficits, they displayed impairments on the WCST similar to those observed in frontal lobe subjects. The deficit was less pronounced in patients taking the DA precursor L-Dopa (Bowen, Kamienny, Burns & Yahr, 1975), which has therapeutic efficacy in re-establishing dopaminergic tone.

In view of these findings, several authors have proposed that reduced dopaminergic tone in prefrontal cortex may be associated with hypofrontality in schizophrenia, and may be responsible for several of the cognitive deficits that have been observed. Levin (1984) has reviewed a wide variety of behavioral data in support of this conjecture (also see Levin, Yurgelun-Todd & Craft, 1989). Weinberger and his collaborators have focused on impairments in WCST performance (Weinberger et al., 1986; Weinberger, Berman & Chase, 1988). In one study (Weinberger, Berman & Illowsky, 1988), this group showed that levels of the DA metabolite HVA in the cerebrospinal fluid of schizophrenic patients showed a very strong correlation with prefrontal activity (as measured by rCBF) during WCST in schizophrenic patients. In another study (Geraud, Arne-Bes, Guell & Bes, 1987), the hypofrontality in schizophrenics observed by PET was reversed with administration of a DA agonist. Thus, there is mounting evidence that DA is closely related to the activity of frontal

cortex, and that a disturbance in this system may be involved in schizophrenic cognitive deficits.

C. Summary

We began by reviewing behavioral data that suggest schizophrenics suffer from the inadequate processing of context. In particular, schizophrenics appear unable to maintain contextual information necessary for the control of action. We then reviewed biological data addressing: a) the role of prefrontal cortex in maintaining context; b) the involvement of the mesocortical dopamine system in the function of prefrontal cortex; and c) disturbances of both of these systems in schizophrenia. Despite a growing recognition that these findings are related, no theory has been proposed yet which explains — in terms of causal mechanisms — the relationship between disturbances in frontal cortex and dopamine on the one hand, and behavioral deficits on the other. In the next section, we introduce a set of information processing mechanisms, and models based on them, that can provide a causal explanation of behavioral deficits in terms of specific disturbances at the biological level.

III. Connectionist Simulations of Biological and Behavioral Disturbances

In this section we will present three information processing models which simulate performance in each of the three experimental tasks we discussed above: the Stroop task, the CPT and the lexical disambiguation task. These models all contain a distinct component for maintaining context that can be identified with the function of frontal cortex. Furthermore, they implement a specific mechanism for the influence of dopamine in frontal cortex. In simulations using these models we show that an impairment of this mechanism results in patterns of performance that are analogous to that of schizophrenic patients in these different tasks. As background for understanding the models, we first provide a brief overview of the framework within which they were developed.

A. The Connectionist Framework

The models draw upon the principles of parallel distributed processing (Rumelhart & McClelland, 1986a; McClelland & Rumelhart, 1986), or connectionism. These principles provide a framework for building computer models that can simulate cognitive phenomena. In particular, these principles are meant to capture the salient details of the mechanisms underlying information processing as it occurs in the brain. By doing so, it is hoped that: a) this will lead to more realistic models of cognitive phenomena; and b) it will be possible to relate behavior directly to biological processes (for an indepth discussion, see Rumelhart & McClelland, 1986b). Connectionist models have been used effectively to explain a variety of phenomena, both at the biological and behavioral levels. These include the computation of spatial orientation from retinal and eye-position information (Zipser & Andersen, 1988), the computation of object shape from shading information (Lehky & Sejnowski, 1988), the acquisition of regular and irregular verb forms in English (Rumelhart & McClelland, 1986c),

text to speech translation and disturbances of this phenomenon in surface dyslexia (Seidenberg & McClelland, 1990), and access to word meaning from word form in deep dyslexia (Hinton & Shallice, 1989). The principles of the connectionist framework can be roughly divided into those having to do with processing, and those having to do with training.

Processing. Each unit in a connectionist network is a simple, summing device: it accumulates inputs from other units, and adjusts its output in response to these inputs. Typically, units are grouped into modules, and modules are connected into pathways. Information is represented as the pattern of activation over the units in a module. The activation of each unit is a real valued number varying continuously between a minimum and maximum value, which can be thought of as the unit's probability of firing. The responsivity of each unit is scaled by its gain parameter, which serves as a multiplier for the effects of excitatory and inhibitory inputs to the unit. Processing occurs by the propagation of signals (spread of activation) among units within and between modules. This occurs via the connections that exist between units. The connections between the units of different modules constitute processing pathways.

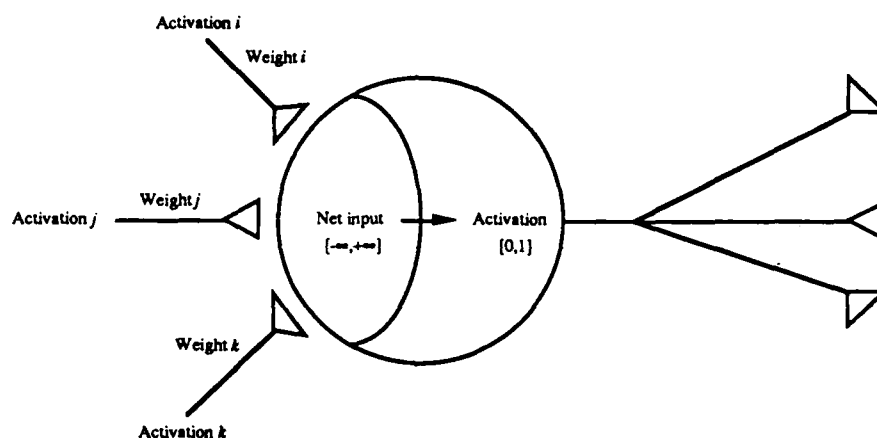


Figure 3: Schematic of a typical unit in a connectionist system

Training. The ability of this type of system to perform a given task depends on its having an appropriate set of connection weights in the pathway that runs from the input module(s) to the output module(s) relevant to the task. The connections in a pathway are set by learning. Although a number of different connectionist learning techniques have been described, the generalized delta rule, or back propagation algorithm (Rumelhart, Hinton & Williams, 1986) is in widest use. In brief, this involves the following series of operations: 1) present an input pattern to the network; 2) allow activation to spread to the output level; 3) compute the difference (error) for each output unit between its current activation and the one desired (i.e., the one specified by the target, or teaching pattern); 4) "back propagate" these error signals toward the input units. The back propagation algorithm provides a way for each unit in a

pathway to compute the adjustment it must make to its connection weights so as to best reduce the error at the output level.⁴

It is important to recognize that connectionist models are not usually meant to be detailed circuit diagrams of actual neural networks. Rather, like statistical mechanical models in physics and chemistry, connectionist models are designed with the intention of capturing the features of a lower level system (information processing mechanisms in the brain) that are most relevant at a higher level of analysis (cognition and behavior). Thus, an important goal of such models is to examine the effects of biological variables on behavior, without having to reproduce the entire brain in order to do so.

Using the connectionist framework, we have developed simulation models of three tasks relevant to research on schizophrenia: the Stroop task, the continuous performance test, and the lexical disambiguation task described above. Each model simulates normal performance in one of these tasks. All three models share a common mechanism for processing context. This relies on a specific module which we identify with frontal lobe function. In each model, reducing the gain of units in this module is sufficient to produce the pattern of performance observed for schizophrenics in the corresponding task. We begin our description of the models by showing how the physiological influence of dopamine can be simulated by changes in the gain parameter of individual units. We then describe simulations of normal and schizophrenic performance in the Stroop, CPT and lexical disambiguation tasks.

B. Simulation of the Physiological effects of Dopamine

In contrast to other neurotransmitter systems such as aminoacids or peptides, the anatomy and physiology of dopamine systems are not suited to the transmission of discrete sensory or motor messages. Rather — like other catecholamine systems — dopamine systems are in a position to modulate the state of information processing in entire brain areas over prolonged periods of time. Several anatomical and physiological observations support this contention. Dopamine neurons originate in discrete nuclei localized in the brain stem and their fibers project radially to several functionally different areas of the CNS. The baseline firing rate of these neurons is low and stable, and the conduction velocity along their fibers is slow. These characteristics result in a steady state of transmitter release and relatively long-lasting post-synaptic effects, that are conducive to modulatory influences. Most importantly, recent evidence suggests that the effect of dopamine release is not to directly increase or reduce the firing frequency of target cells (e.g., Rolls, Thorpe, Boytim, Szabo & Perrett, 1984; Chiodo & Berger, 1986). Rather, like norepinephrine, dopamine modulates the response properties of post-synaptic cells such that both inhibitory and excitatory responses to other afferent inputs are potentiated. Some investigators have described this effect as an increase in the 'signal-to-noise ratio' of the cells' behavior (Foote, Freedman & Oliver, 1975) or an 'enabling' of the cells' response (Bloom, Schulman & Koob, 1989).

⁴ A common criticism of this algorithm is that it is not biologically plausible. That is, it is difficult to imagine that real neural systems rely on the back propagation of error signals for learning. However, back propagation implements the general phenomenon of "gradient descent" — the gradual reduction of error by incremental adjustments in connection weights. Gradient descent has proven to be a powerful concept for describing many of the details concerning human learning behavior. Thus, it may be that back propagation offers a reasonable approximation of the type of learning that occurs in neural systems, even if the actual algorithm is different.

Experiments investigating the modulatory effects of norepinephrine have been performed in many regions of the brain. However, the modulatory effects of dopamine have been investigated mostly in the striatum (e.g., Rolls et al., 1984; Chiodo and Berger, 1986). In such studies, the modulatory effects of dopamine — mediated by either D1 or D2 receptors (Hu and Wang, 1988) — are similar to those observed for norepinephrine. However, researchers have only recently begun to look at the effects of dopamine in prefrontal cortex. Two studies (Aou et al., 1983; Sawaguchi & Matsumura, 1985) reports inhibitory as well as excitatory effects of dopamine iontophoresis in primate prefrontal cortex, as a potentiation model would lead us to expect. Other studies (Reader, Ferron, Descarries & Jasper, 1979; Ferron, Thierry, Le Douarin & Glowinski, 1984; Sesack & Bunney, unpublished observations) report a potentiation of inhibitory responses but a *reduction* of excitatory responses. This contrasts with the potentiation of excitatory as well as inhibitory responses that we would expect, and that has been observed in striatal cells. However, in these latter studies the amount of dopamine released also produced a direct decrease in the baseline firing rate of the cells. This direct decrease in baseline firing rate has been observed in striatal cells but only when large amounts of dopamine were released, and not for smaller amounts (Chiodo & Berger, 1986).⁵ Thus, the reduction of excitatory responses observed in these studies may be related to the use of high concentrations of dopamine. The effects of smaller concentrations — which do not affect baseline firing rate — have not been tested in the prefrontal cortex.

In our models, we assume that the effects of dopamine on cells in prefrontal cortex — at concentrations relevant to the behavioral tasks we are interested in — are similar to the effects that have been observed in striatal cells: a potentiation of inhibitory as well as excitatory inputs. This assumption may turn out to be too simplistic or even incorrect. Nevertheless, an important function of our models is to establish a dialogue between biological and behavioral research, from which both can benefit. By assuming that dopamine has the same effects in prefrontal cortex as it does in the striatum, we have begun to account — in the simulations described below — for a variety of behavioral findings associated with schizophrenia. Success of our models at the behavioral level can be taken as a prediction concerning the validity of our assumptions at the biological level. In this way, the models provide motivation and theoretical guidance for further studies at the physiological level concerning the effects of dopamine in prefrontal cortex.

For the purposes of simulating the potentiating effects of dopamine, we assume that the relationship between the strength of the afferent input (excitatory or inhibitory) to a neuron and its frequency of firing can be represented, in a connectionist network, as a non-linear function relating the net input of a unit to its activation value. Physiological experiments suggest that in biological systems the shape of this function is sigmoid, with its steepest slope around the baseline firing rate (e.g., Freeman, 1979; Burnod & Korn, 1989). The same experiments also indicate that small increments in excitatory drive result in greater changes in firing frequency than equivalent increments in inhibitory input. These properties can be captured by the logistic function with a constant negative bias:

⁵ Indeed, the concentration of dopamine in the iontophoresis micropipettes and the intensity of the ejection current used in Reader et al. (1984) were both one order of magnitude greater than the concentrations and current intensity used in Chiodo & Berger (1986). In Reader et al. (1984) the cells responded to these levels of dopamine release with an almost complete cessation of spontaneous firing whereas in Chiodo & Berger (1986) the baseline firing rate was not affected.

$$\text{activation} = \frac{1}{1 + e^{-(\text{gain} * \text{netinput}) + \text{bias}}} \quad \text{Equation (1)}$$

(also see Figure 4, Gain = 1.0). Furthermore, the effects of dopamine — potentiation of unit response — can be simulated by increasing the gain parameter of the logistic function. As Figure 4 (Gain = 2.0) illustrates, with a higher gain the unit is more sensitive to afferent inputs while its baseline firing rate (net input = 0) remains the same. We have shown that such a change in gain can simulate a number of different catecholaminergic effects at both the biological and behavioral levels (e.g., the influence of catecholamines on the receptive field of individual units, the influence of amphetamines on stimulus detection, and stimulus response generalization in both humans and rats — Servan-Schreiber, 1989).




Figure 4. The influence of the gain parameter on the logistic activation function of an individual unit. Note that, with an increase in gain, the effect of the net input on the unit's activation is increased, while the reverse is true with a decrease in the gain. These effects simulate the consequences of dopamine release on target neurons in the CNS.

As we noted earlier, the anatomy of catecholamine systems suggests that they can influence whole regions of the CNS at once. In particular, the mesocortical dopamine system implicated in schizophrenia has extensive projections to frontal cortex. To model the action of dopamine in this region of the brain, we simply change the gain parameter of all the units in the module supporting the function corresponding to this region. In the models we report below, decreased dopaminergic supply to frontal cortex was simulated by reducing the gain of units in the module used to represent and maintain context. In all three models, simulation of schizophrenic performance was conducted by reducing gain from a normal value of 1.0 to a lower value in the range 0.6-0.7.

C. A connectionist Model of Selective Attention (the Stroop Effect)

Architecture and processing. Elsewhere, we have described in detail a connectionist model of selective attention that simulates human performance in the Stroop task (Cohen, Dunbar & McClelland, 1990). In brief, this model consists of two processing pathways, one for color naming and one for word reading (see Figure 5). Simulations are conducted by activating input units corresponding to stimuli used in an actual experiment (e.g., the input unit representing the color red in the color naming pathway) and allowing activation to spread through the network. This leads to activation of the output unit corresponding to the appropriate response (e.g., "red").




Figure 5. Network architecture. Units at the bottom are input units, and units at the top are the output (response) units.

Training and attentional modulation. The model is trained to produce the appropriate behavior by presenting it with the input patterns for each of the responses it is expected to make, and using the back propagation learning algorithm to adjust the connection weights

accordingly. During training, the model is given more experience with (i.e., a greater number of training trials on) the word reading task than the color naming task,⁶ so that the connection weights in the word reading pathway become greater than those in the color naming pathway. As a result, when the network is presented with conflicting inputs in the two pathways (e.g., the word RED with the color green), it responds preferentially to the word input. In order to modulate this effect, the system is equipped with a set of units that are used to represent the intended behavior (i.e., color naming vs. word reading). Thus, task specification is represented by the appropriate pattern of activation over a set of "task demand" units. These task demand units connect to the intermediate units in each of the two pathways and, based on the pattern of activation over the task demand units, modulate the responsivity of the units in the two processing pathways. For example, when the pattern corresponding to "color naming" is activated over the context units, activation spreading from these units has a potentiating effect on processing units in the color pathway, while it "desensitizes" units in the word pathway. This produces a modulation of the flow of information in the two pathways, favoring the color pathway. The result is that, although the connection strengths in the color pathway are weaker, a signal presented to this pathway is able to overcome the dominant response otherwise mediated by the word pathway. In this way, the model is able to selectively attend to information in the task-relevant pathway.

Simulation. This simple model is able to simulate an impressive number of empirical phenomena associated with the Stroop task. It captures all of the phenomena depicted in Figure 1 (asymmetry in speed of processing between word reading and color naming, the immunity of word reading to the effects of color, and the susceptibility of color naming to interference and facilitation from words, and greater interference than facilitation), as well as the influence of practice on interference and facilitation effects, the relative nature of these effects, response set effects and stimulus onset asynchrony effects (see Cohen, Dunbar & McClelland, 1990).

This model also exhibits behaviors that make it relevant to understanding schizophrenic disturbances of attention. In the model, the task demand units — which are responsible for attentional selection — can be thought of as maintaining the context necessary for selection of the task-relevant response. When subjects are presented with conflicting input in two dimensions (e.g., the word GREEN in red ink), they respond to one dimension and not the other, depending upon the *context* in which it appears (i.e., the task: color naming or word reading). If frontal cortex is responsible for maintaining this context (i.e., the task demand representation), and if schizophrenia involves a disturbance of frontal lobe function, then we should be able to simulate schizophrenic performance in the Stroop task by disturbing processing in the task demand module. More specifically, if frontal lobe dysfunction in schizophrenia is due to a reduction in the activity of its dopaminergic supply, then we should be able to simulate this by reducing the gain of units in the task demand module.

Figure 6 shows the results of such a simulation, in which the gain of only the task units was reduced; all other units were unperturbed. This change in the context (task demand) module produces effects similar to those observed for schizophrenics: an increase in overall response time, with a disproportionate increase for color naming interference trials. Thus, the model shows that a lesion restricted to the mechanism for processing context can produce both an

⁶ This corresponds to the common assumption that human adults have had more experience generating a verbal response to written words than to colors they see.

overall degradation in performance as well as the expected attentional deficit (i.e., in the interference condition).

The model also allows us to compare the effects of this specific disturbance to those of a more general disturbance, addressing a common difficulty in schizophrenia research. Recall the argument that, in the context of a general degradation of performance in schizophrenics (e.g., overall slowing of response), it is difficult to know whether degradation in a particular experimental condition is due to a specific deficit or a more generalized one. However, this difficulty arises primarily when the mechanisms for the deficits involved have not been specified. The model provides us with a tool for doing this. Above, we described the mechanism for a specific attentional deficit. To compare this to a more generalized deficit, we induced overall slowing in the model by decreasing the rate at which information accumulated for each unit (cascade rate); this was done for *all* of the units in the model. The results of this manipulation appear in Figure 6. When the cascade rate is decreased, there is an overall slowing of response, but no disproportionate slowing in the interference condition. In contrast, as we noted above, the specific attentional disturbance produces both effects: slowing occurs in all conditions, but this is most pronounced in the interference condition. Thus, the attentional hypothesis provides a better account for the data than at least one type of generalized deficit. We have explored others (e.g., an increase in the response threshold), with similar results.




Figure 6. Stroop task performance for normal and schizophrenic subjects, and results from simulations manipulating the gain parameter (task demand units only) and cascade rate (all units) in the network. For the empirical data, response times are the number of seconds to complete each card in the classic version of the Stroop task; for the simulations, response times are the average number of processing cycles required to respond to the stimuli in each condition.

The model we have described relates a disturbance in attention directly to the processing of context. Attention is viewed as the effects that context has on processing, and a failure to maintain an appropriate contextual representation (e.g., the task demand specification) leads directly to a failure in selective attention. In the Stroop task, this manifests as an increased susceptibility to interference in the color naming task. This, in turn, reflects the increased

influence of dominant response processes (e.g., word reading) that occurs with the weakening of attention. Schizophrenic performance has often been characterized as reflecting a dominant response tendency (e.g., Chapman, et al., 1964; Maher, 1972), although no specific mechanism has previously been proposed for this. We will return to this issue in our discussion of schizophrenic language performance below.

D. Simulation of the Continuous Performance Test (CPT)

The Stroop model shows how contextual information and its attentional effects can be represented in a connectionist model, and how a disturbance in this mechanism can explain schizophrenic performance deficits. The principles demonstrated by this model have general applicability. In this section we discuss their extension to a task in which *memory* for context is more directly involved.

In the Stroop task, the context consists of instructions to attend to ink color or to words. These instructions remain constant throughout the task, so they are not difficult to maintain. However, in other attentional tasks the context necessary to select a response is derived from previous stimuli. For example, in a version of the CPT called CPT-double, targets consist of any consecutively re-occurring letters (e.g., 'B' immediately following a 'B'). In this case, memory for context (the previous stimulus) is necessary to select the appropriate response. Without this, each letter would be ambiguous: it could be a target or a distractor. Therefore, an impairment in memory for context should degrade performance, such as is observed for schizophrenics. To demonstrate this, we constructed a network to perform the CPT-double task.




Figure 7. Network used to simulate the CPT-double. Note the bidirectional connections from units in the input and letter modules.

Architecture, processing and training. The network consisted of four modules: an input module, an intermediate (associative) module, a letter module and a response module (see Figure 7). The input module was used to represent visual features of individual letters as they may appear on a computer screen. Different letters consist of different features so that each one produced a different pattern of activation over the input units. The network was trained to associate these activations patterns with their corresponding letters, by activating the appropriate units in the letter identification module. In addition, the network was trained to activate the response module unit whenever a letter appeared a second time in a row. To do this, however, the network must be able to store and use information about the previous as well as the current stimulus. We made this possible by introducing a set of connections from the letter units back to the intermediate units. Thus, intermediate units received "bottom up" information from the feature units (representing the current input) and "top down" information from the letter units (representing the network's interpretation of the previous input). In this way, the network could compare the current and previous inputs, and learn to activate the response unit when two consecutive letters were identical. Note that there is a direct analogy between the role played by the letter units in this model and the role played the task demand units in the Stroop model. That is, the representation over the letter units in the CPT model provided the context for disambiguating the response to a particular pattern of input just as the task demand units did in the Stroop model. In the CPT model, however, context was determined by the previous input, and therefore changed from trial to trial.

Simulation. Following training, the network was able to perform the CPT-double task perfectly for a set of 26 different stimuli. To simulate the performance of normal subjects — who typically show errors of omission on approximately 13% of trials, and errors of commission on approximately 1% of trials (see Figure 8A) — we added noise to processing. Noise in neural systems is usually attributed to sources of afferent input that are independent of the relevant stimulus. To simulate this distortion of the input to a unit, we added a small amount of random, normally-distributed noise to the net input of every unit on each processing cycle. The amount of noise (standard deviation of the distribution) was adjusted to match the performance of the network with that of human subjects. The results of this simulation appear in Figure 8B (gain = 1.0). Then, to simulate schizophrenic performance, we disturbed processing in the letter module, which was responsible for maintaining context. As in the Stroop simulation, we decreased the gain of these units and observed its behavior. The percentage of misses increased to 20% and false alarms increased slightly to 1.1%. These numbers closely match the results of empirical observations of schizophrenic subjects (see Figure 8). Although some have interpreted schizophrenic CPT performance in terms of a deficit in sensory processing, our model suggests an alternative hypothesis: performance deficits are due to a degradation in the memory trace required — as context — for processing the current stimulus. This hypothesis is consistent with our account of Stroop performance, and with disturbances of language processing that we turn to next.




Figure 8. Percentage of misses and false alarms for normal and schizophrenic subjects in the CPT task (panel A), and for the simulation run with normal and low gain on units in the Letter Identification Module (panel B).

E. Simulation of Context-Dependent Lexical Disambiguation

The previous simulations show how deficits in two tasks, which on the surface appear to be very different, can be understood in terms of a common set of processing mechanisms. The Stroop model showed how an overall increase in reaction time and a dominant response bias can result from a poor representation of context. However, this task did not directly address the role of memory for context in processing. This was addressed by the CPT model. In this case, however, no dominant response bias was apparent because the task did not involve any dominant response tendencies. The lexical disambiguation task we described earlier provides an opportunity to examine both dominant response bias and poor memory for context at once. The results of our study replicated the findings of others that schizophrenics show a tendency to respond to the dominant meaning of lexical ambiguities, even when context confers the weaker, less frequent meaning. However, our results suggested that this tendency is significant only when context is temporally remote, implicating a deficit in *memory* for context. We were able to simulate these language deficits using the same mechanisms that were used to account for schizophrenic performance in the Stroop task and CPT.

Architecture and processing. The model used to simulate performance in the lexical disambiguation task (Figure 9) employed the same basic architecture as the CPT model (see Figure 7). The input module was used to represent lexical stimuli (e.g., the word PEN). The network was trained to associate patterns of activation in this module with patterns in two of the other modules: the output module and the discourse module. Patterns in the output module represented an overt response to the meaning of the input word (e.g., "writing implement"), while the discourse module represented the topic of the current sequence of inputs (e.g., the meaning of the sentence or phrase, rather than the meaning of individual words). The

intermediate module functioned as a semantic module, encoding an internal representation for the meaning of the input that was used to generate an appropriate response in the output module, and a relevant discourse representation in the discourse module. As in the CPT model, there were two-way connections between the semantic module and the discourse module. This meant that not only could an input generate a discourse representation but, conversely, once a discourse representation had been activated, it could have a "top down" influence on processing in the semantic module. This provided the mechanism by which context could be used to resolve lexical ambiguity.

Training. The model was trained to produce an output and discourse representation for 30 different input words, some of which were ambiguous. For ambiguous words, the network was sometimes trained to produce the response and discourse pattern corresponding to one meaning of the ambiguity (e.g., PEN → "writing implement" and *WRITING*),⁷ while on other trials it was trained to produce patterns corresponding to the other meaning (e.g., PEN → "fenced enclosure" and *FARMING*). The network was trained more on one meaning than the other. This asymmetry of training was similar to that of the Stroop model (trained on words more than colors), with a comparable result: when presented with an ambiguous input word, the network preferentially activated the dominant (more frequently trained) response and discourse representations. To permit access to the weaker meaning, the network was sometimes presented with an ambiguous word along with one of its associated discourse representations as input (e.g., PEN and *FARMING*),⁸ and trained to generate the appropriate response (i.e., "fenced enclosure"). Finally, the network was trained on a set of context words, each of which was related to one meaning an ambiguity. These words (e.g., CHICKEN) were trained to produce their own meaning as the response ("fowl"), and a discourse representation that was the same as for the related meaning of the ambiguity (*FARMING*).

⁷ We will refer to input words in upper case (no italics), to output responses in quotation marks, and to discourse representations in italicized upper case.

⁸ Recall that the discourse module is connected to the semantic module with two way connections, so that the discourse module can be thought of as either an input module or an output module, depending upon whether the representation in this module is explicitly specified by the experimenter, or is allowed to develop in response to activation it receives from the semantic module.




Figure 9. Schematic diagram of the language processing model. Patterns of activation over the units in the input module are assumed to represent the current sensory stimulus (e.g., the orthographic code for a written word), while the output module is assumed to represent the information necessary to generate an overt response (e.g., the phonological code needed to pronounce the meaning of the word). Note that the connections between the semantic and discourse modules are bidirectional.

The combined effects of these training procedures was that when an ambiguous word was presented and there was no representation active over the discourse units, the output was a blend of the two meanings of the word, with elements of the more frequently trained (dominant) meaning being more active than the other (subordinate) meaning. However, when a discourse representation was active, the model successfully disambiguated the input and activated only the relevant output.

Simulation. Trained in this way, the model was able to simulate the use of context in natural language processing. Most words in English have more than one meaning. Furthermore, language is sequential. Therefore, processing language requires memory of the context provided by prior stimuli to disambiguate current ones. In the model, this occurs by constructing a discourse representation in response to each lexical input that can then be used as context for processing subsequent stimuli. We tested the model for this ability by presenting it with a word related to one of the meanings of an ambiguity (e.g., CHICKEN), then presenting the ambiguity (e.g., PEN) and observing the output. Note that in this case, the model was not directly provided with a discourse representation. Rather, it had to construct this from the first input, and then use it to disambiguate the second. Tested in this way with all context-word/ambiguous-word pairs (e.g., either CHICKEN or PAPER followed by PEN), the model was consistently able to generate the context-relevant meaning response to each ambiguity.

To simulate performance in our lexical disambiguation experiment, the model was presented with pairs of context and ambiguous words (representing the clauses used in the experiment) in either order (context word first or last). Following each pair, the network was probed with the ambiguous word, simulating the subjects' process of reminding themselves of the ambiguity, and choosing its meaning. At each time step of processing, a small amount of noise was added to the activation of every unit in the model. The amount of noise was adjusted so that the simulation produced an overall error rate comparable to that observed for control subjects in the experiment. The model's response in each trial was considered to be the meaning that was most active over the output units after the probe was presented. To simulate schizophrenic performance, we introduced a disturbance analogous to the one used in the Stroop and CPT models: a reduction in gain of units in the module responsible for maintaining context. In the current model, this was the discourse module. The results of these simulations show a strong resemblance to the empirical data (see Figure 10). They demonstrate both effects: a) in the low gain mode, the simulation made about as many more dominant response errors as did schizophrenic subjects; however, b) as with human subjects, this only occurred when context came first. The simulation also captured the reverse trend observed among control subjects: fewer errors when context came first. The number of unrelated errors made by the model (not shown in Figure 10) was approximately the same in both the low gain and normal gain mode, as was the case across groups in the empirical study.

The model provides a clear view of the relationship between dominant response bias and memory for context. When gain is reduced in the context module, the representation of context is degraded; as a consequence, it is more susceptible to the cumulative effects of noise. If a contextual representation is used quickly, these effects are less, and the representation is sufficient to overcome a dominant response bias. However, if time passes (as when context is presented first), the effects of noise accumulate, and the representation is no longer strong enough to mediate the weaker of two competing responses.⁹

⁹ It is worth noting that, when gain is normal in the discourse module, the cumulative effects of noise are offset by a priming effect. That is, when the context representation is sufficiently strong, then its occurrence before the ambiguity allows it to prime the correct meaning, leading to better performance than when context occurs after the ambiguity. Interestingly, a trend toward this effect can also be observed in the empirical data for the control subjects.

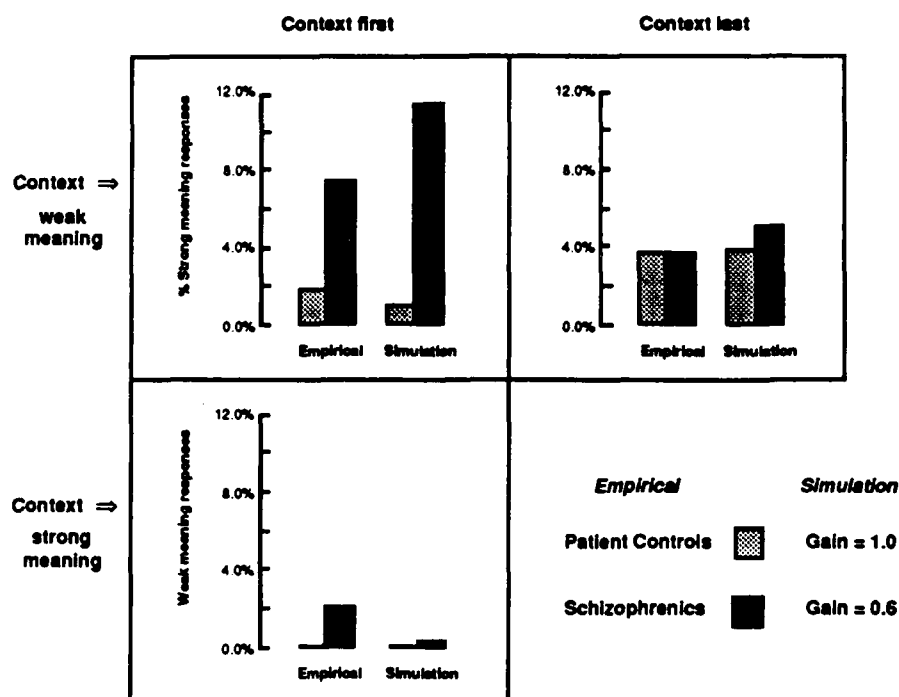


Figure 10. Error rates for subjects in each of the three conditions run in the empirical study and for the the language model simulation. The rate of unrelated errors (not shown) were the same in the normal and low gain conditions of the simulation, and the same magnitude as that observed in human subjects (about 1-2%).

As we have noted, there are important similarities between this model of language processing and the models of the attentional tasks described earlier. All of the models use a context representation to mediate a response to an ambiguous input. For the Stroop task, the context representation was the pattern of activation over the task demand units; for the CPT it was the pattern over the letter identification units; and for lexical disambiguation it was the pattern in the discourse module. In both the language model and the Stroop model, context was necessary to mediate a weaker response in the presence of a competing, dominant response. In the Stroop model we talked about this as an attentional effect, and deficits as an increase in interference; in language processing it is more common to refer to context effects, and deficits in terms of a dominant response bias. However, our simulation results suggest that the same principles can account for both sets of phenomena. In particular, attention can be thought of as the ability to use context to produce task-relevant behavior; a failure to do so results in the prevalence of dominant response tendencies or, when these do not exist (as in the CPT), a non-specific degradation of performance. Finally, in the CPT and language models, *memory* for context was particularly important. Once again, similar mechanisms were used to in each case, and were able to simulate both normal and abnormal patterns of performance in these very different tasks.

IV. General Discussion

We began by reviewing evidence concerning deficits of attention and language processing in schizophrenia. We also reviewed data which indicate that prefrontal cortex and its dopaminergic supply are important for the processing of context, and that a disturbance in this system is involved in schizophrenia. We then showed how the connectionist framework can be used to relate these findings to one another. We presented three models that: a) simulated quantitative aspects of performance in the Stroop task, a standard version of the CPT, and a lexical disambiguation task; b) elucidated the role of memory for context in both the attentional and language processing tasks; c) related behavior in these tasks to biological processes; and d) identified a specific disturbance in these processes that could account for schizophrenic patterns of performance. The models touch on a number of important issues concerning cognition in both normal subjects and schizophrenics, and the biological processes involved. We discuss these below, as well as some of the limitations of our models. We then compare our models with others which address similar issues. We conclude with a discussion of some general issues concerning the modelling endeavor itself.

1. Attention and Context

The Stroop task and CPT are commonly thought of as measures of attention, whereas the lexical disambiguation task is most naturally thought of as a measure of context effects in language processing. Our models suggest, however, that there is a close relationship between attention and the processing of context. The attentional effects observed in our simulations of the Stroop task and CPT resulted directly from the influence of context. In both cases, the use of context led to the selection of the appropriate response to an otherwise ambiguous stimulus. Similar processes were at work in the lexical disambiguation task, in which context was also necessary for the selection of an appropriate response. This similarity between the attentional and language tasks was demonstrated by our ability to simulate performance in these different tasks using the same basic mechanisms for representing and processing context in each case. Thus, the models contribute to our understanding of the cognitive processes involved in these tasks in two important ways: 1) The models suggest that attention can be thought of as the influence that context has on the selection of task appropriate information for processing, and they are explicit about the mechanisms by which this occurs. 2) The similarity of these mechanisms across tasks suggests that, while at the surface they may appear to be very different from one another, they are governed by a common set of information processing principles. This should not be taken to suggest, however, that the actual processing pathways are the same for all three tasks. Each involves a different level of information processing (from letter recognition to the access of semantic and discourse level knowledge). No doubt, the processing pathways involved at each level are different in ways not captured by our current models. However, these differences do not appear to be relevant to the dimensions of performance we have addressed. Indeed, it is precisely the simplifications introduced by the models that helped bring the commonalities among these tasks into focus.

2. Disturbance in the Processing of Context

Viewing attention as the effects of context also helps organize several findings in the schizophrenia literature. We were able to show that a single disturbance in the mechanisms underlying the processing of context can account for a number of attention and language deficits in schizophrenics — phenomena that have often been treated as separate in the literature. From this, we would predict that performance should correlate across tasks which rely on the processing of context.

Previous attempts to examine cross-task correlations of schizophrenic cognitive deficits have produced mixed results. Kopfstein and Neale (1972) reported small correlations between five different tasks that were presumed to tap attentional mechanisms (a reaction time task, size estimation, the Benjamin proverbs test, the Goldstein-Scheerer object sorting task, and an auditory discrimination task); Asarnow and MacCrimmon (1978) found that performance on the simple CPT-X did not correlate with performance on the span of apprehension test (SAT); while Kornetsky and Orzack (1978) found that the poorer performers on the CPT-X were also more affected by irrelevant preparatory intervals on the Shakow reaction time task. The disparate nature of these findings has led investigators to assume that "attention" may not be a general mechanism. Thus, the measures used in these studies may have tapped different components of attention, or other information processing mechanisms altogether. This may reflect one of the major difficulties that has been faced by this area of research: the lack of a theoretical framework within which to compare and select tasks. Our models offer an approach to this problem, by specifying the mechanisms underlying at least one component of attention (the effects of context), and relating these directly to task performance. In particular, they identify two task dimensions that are relevant to attentional effects, and schizophrenic deficits: a) the relative strength of competing responses and b) the demands placed on memory for context. Table 2 categorizes the tasks we have considered along these dimensions.

Table 2. Memory for Context and Response Strength

	<i>Less memory for context</i>	<i>More memory for context</i>
<i>Equivalent response strengths</i>	CPT-X	CPT-AX CPT-Double
<i>Asymmetric response strengths</i>	Stroop task (interference condition)	Lexical disambiguation task (context first)

Tasks in which subjects need to keep only a set of instructions or a single stimulus in mind place the least demand on memory for context. That is, when task instructions or a target stimulus remain constant throughout the task, they are reinforced by performance on each trial, and therefore rely less on memory. For example, in the CPT-X (detect any occurrence of an "X") the subject needs to remember only the target stimulus, and in the standard Stroop paradigm — in which trials are blocked by task — the instructions remain constant (respond to color or respond to word). These tasks are shown on the left side of Table 2. In contrast, in the CPT-double and CPT-AX subjects must remember the previous stimulus in addition to the task instructions, increasing the demand placed on memory for context. This is also true of the

lexical disambiguation task, when context comes first. These are shown on the right side of Table 2. Attentional effects related to the processing of context should be most evident in these tasks.

The second dimension of Table 2 concerns the influence of competing response tendencies. In some tasks, all potential responses are of equivalent strength. For example, subjects are equally familiar with the letters used in standard variants of the CPT; these are shown at the top of Table 2. In other tasks, however, the strength of one response is much greater than the strength of the other. This is due to different amounts of experience either with different aspects of the stimulus (as in the Stroop task: colors vs. words), or with different responses to the same stimulus (as in the lexical disambiguation task). In our simulations, these differences were captured by differential amounts of training on competing stimulus-response associations. Tasks with response strength asymmetries are shown at the bottom of Table 2. While contextual effects can be observed whenever a stimulus is associated with more than one response, tasks in which competing responses are of unequal strength will be most sensitive to these effects. For example, in the CPT-double failure to use context would result in performance at chance (this is because, in the absence of context, the competing responses have equal strength). A much stronger effect would be observed in the Stroop and lexical disambiguation tasks: a consistent elicitation of the stronger response, even when it is inappropriate. Thus, the latter should provide the most sensitive measure of attentional effects related to the processing of context.

Table 2 bears directly on schizophrenic deficits in these various tasks. To the extent that a disturbance in the processing of context is involved, we would expect performance to be least affected in tasks at the top and to the left of Table 2. Existing data support at least one implication of this analysis: schizophrenics show fewer and less reliable deficits in the CPT-X than the CPT-double or CPT-AX (Nuechterlein, 1984). We may also be able to explain one of the failures to correlate across measures of attention. Asarnow and MacCrimmon (1978) found no relationship between performance in the CPT-X and SAT. As in the CPT-X, target stimuli in the SAT remain constant throughout the task; this task belongs with CPT-X in the upper left of Table 2. Since these tasks are the least sensitive to context, we would expect them to be the least likely to reveal a correlation.

Most importantly, Table 2 provides a rational approach for the design of new studies to evaluate cross-task correlations. Schizophrenics should show the greatest deficits, and therefore the greatest correlations, when the tasks involve both memory for context and a response strength asymmetry (i.e., dominant response tendency). We have begun to provide support for this prediction with the results of our lexical disambiguation task. However, it should be possible to demonstrate increased sensitivity to schizophrenic deficits in each of the other tasks — and correlations among them — by varying them along the appropriate dimensions of Table 2. For example, response strength asymmetry could be introduced into the CPT by varying the frequency with which targets appear. Conversely, the reliance on memory for temporary context information could be increased in the Stroop task by varying the instructions from trial to trial, and presenting the stimuli at various delays following the instructions. These task manipulations should increase both their sensitivity to schizophrenic deficits and the likelihood of detecting cross-task correlations.

We should be clear, of course, that a disturbance in the processing of context may be only one of several disturbances underlying schizophrenic cognition. Indeed, we have focussed on a circumscribed set of experimental findings in this paper. From a clinical perspective, these may represent cognitive correlates of the "negative symptoms" of schizophrenia: flattening of affect, emotional withdrawal and amotivation. While both may be related to frontal lobe

deficits, the models in their present form do not provide an account of the relationship between the cognitive and affective manifestations of this disorder. We also have not addressed the "positive symptoms" of schizophrenia, such as hallucinations and delusions. It is possible, however, that the mechanisms we have discussed may be relevant to these symptoms. For example, an *increase* in the gain parameter throughout the network (corresponding to an increase in dopaminergic activity) results in active and contrasted patterns on the output layer, regardless of the strength (i.e., degree of activation) of the input. When such output patterns are produced in the absence of a meaningful input, the network might be considered to display misperceptions or misinterpretations that resemble the phenomena of hallucinations and delusions. A related argument has been offered by Hoffman (1987), which we will discuss below.

3. Generalized versus Specific Deficits

A common issue in schizophrenia research is the extent to which a particular set of findings reflect a generalized deficit as opposed to a deficit in a specific component of processing. For example, the ubiquitous finding of an increase in reaction time is typically considered to reflect a generalized deficit. However, it is difficult to know the meaning of this hypothesis without defining it in specific information processing terms. A generalized deficit must still reflect a disturbance of *some* kind, *somewhere* in the system. Our model of the Stroop effect provided one possible interpretation of this hypothesis (a slowing in the rate of information processing in all components of the system), and allowed us to compare it with a hypothesis concerning a more specific deficit (a reduction of gain in the module responsible for processing of context). Thus, the model not only provided a framework within which to make these hypotheses explicit, but also to compare their ability to provide quantitative fits to the data. In the case of these two hypotheses, our findings favored the more specific deficit. One implication of this was that what appeared to be a general effect (overall slowing of response) could be attributed to a circumscribed disturbance. While the increase in reaction time for schizophrenics in other tasks may well be due to more general deficits, the Stroop model showed that this need not be the case.

4. Frontal Cortex, Dopamine and the Processing of Context

Fuster (1980) and Goldman-Rakic (1985) have pointed out the role of prefrontal cortex in relating information over space and time. Diamond (1988) has emphasized the importance of the role that this area plays in inhibiting "prepotent" (dominant) response tendencies. Our models show how these information processing functions may be implemented in biologically plausible mechanisms, and how they may be modulated by dopaminergic activity. For example, in our simulation of the lexical disambiguation task, the discourse module supported a representation that was built up in the course of processing. This provided a form of memory that allowed the model to process later elements of the sentence in the context of ones it had seen earlier. The models also showed how such contextual information permitted the expression of a weaker response in the presence of a stronger (more dominant) one. Thus, we were able to account for two important functions that have been attributed to prefrontal cortex in terms of a specific component in our models. Moreover, they suggested an explicit mechanism for dopaminergic effects in prefrontal cortex. By maintaining or increasing the gain of neurons in this area, dopamine may help augment contextual representations against a background of noise. This, in turn, would lead to better preservation of contextual information over time, and more effective control over dominant response tendencies.

Because the models make the relation between dopamine, frontal cortex and the processing of context explicit, predictions can be made about the interplay between these factors that can be tested empirically. We focus here on how dopamine agonists may affect frontal activity, the relation between frontal activity and the processing of context in the CPT, and the relation between frontal lobe function and language processing.

Dopamine agonists and prefrontal activity. Mesocortical projections form a major component of the dopamine system. From this, it might be expected that dopamine agonists would have the general effect, in normal subjects, of increasing metabolic activity in prefrontal cortex. Our models make a somewhat different prediction: To the extent that a task does not rely heavily on memory for context, and involves a set of routine responses, we would predict that the administration of dopamine agonists would not have any effect on the metabolism of prefrontal cortex during the performance of such a task. This is because our models specify that the effect of dopamine release is to potentiate the response of target cells to afferent signals. In tasks which do not rely heavily on the processing context, we assume that there are fewer signals arriving in prefrontal areas, and therefore the activity of units in these areas should be relatively unchanged. However, during performance of a task that *does* rely on memory for context, the effect of dopamine agonists should be to substantially increase metabolism in prefrontal cortex. Thus, we predict an interaction between task-type and drug. Our predictions receive preliminary support from data reported by Geraud et al. (1987). In this study, the prefrontal activity of normal subjects at rest was not increased by the dopamine agonist Pirebdil. However, the second part of our prediction has not yet been tested: that an agent such as Pirebdil would increase prefrontal activity in the same subjects during a task requiring memory for context.¹⁰

Prefrontal activity during the CPT. We have argued that schizophrenic deficits on the CPT can be attributed to frontal lobe dysfunction. In some reports, however, performance on the CPT has failed to differentiate between schizophrenic subjects and controls. Berman et al. (1986) have reported the absence of any correlation between CPT performance and prefrontal activation in either schizophrenics or normal controls; neither group showed significant prefrontal enhancement during the task. These results are in conflict with the findings of R. M. Cohen et al. (described above), in which a significant correlation between prefrontal metabolism and CPT performance was reported. The analysis of task dimensions relevant to CPT performance that we presented above may provide a reconciliation of these findings. Berman et al. used two variants of the CPT: the simple CPT-X which makes fewer demands on memory for context, and a version of the CPT-AX with interstimulus intervals (ISIs) of 0.8 seconds or less. At such short ISIs, the association between A and X can be encoded through direct reinforcement; we assume that reinforcement learning does not rely on prefrontal areas (see discussion above: "Function of prefrontal cortex"). In fact, as subjects' performance improved, Berman et al. attempted to increase the difficulty of the CPT-AX by reducing the ISI even further. Such an increase in event-rate of the task has been shown to impair performance in normal subjects (Parasuraman, 1979). According to our analysis, this increase in difficulty is unrelated to the specific difficulty that the CPT represents for schizophrenics. Rather, it is

¹⁰ Of course, an adequate test of the hypothesis requires that prefrontal metabolism during a context-requiring task be compared with that during a control task that has been matched on all other dimensions, such as intrinsic difficulty, stimulus and response modality, etc.

when the duration between the contextual cue (here the letter 'A') and the potential target ('X') is increased that we would expect them to show difficulty with the task; that is, as their memory for context fails to reliably bridge the gap between the two stimuli. These observations may explain why no specific increase in prefrontal activity was observed during performance of the two variants of the CPT used by Berman et al.

In contrast, R. M. Cohen et al. used an auditory CPT in which subjects were asked to detect the softest of three tones of equal frequency. Tones were presented at two-second intervals. In this task, the target can be identified only in reference to the distractors, which, therefore, have to be actively maintained in memory. Because the subjects need to integrate over several previous trials (at least two) in order to make the relevant comparison, and because of the longer ISI (two seconds), this task places greater demands on memory for context. It is not surprising, therefore, that in this study a correlation was found between prefrontal activity and CPT performance. Based on these arguments, we can make the following prediction: the ability of standard versions of the CPT-AX and CPT-double to differentiate schizophrenic from normals should depend on ISI. When the ISI is of one second or less, schizophrenic performance should not be dramatically impaired compared to normals. However, at longer ISIs (e.g., five seconds) normal subjects should do better (because the event-rate goes down), whereas schizophrenics' performance should be degraded (because memory for context is now required). Moreover, in normal subjects, CPT performance may not correlate with PFC activity at short ISIs (as Berman et al. found), but it should correlate during blocks of trials at longer ISIs.

Prefrontal cortex and language performance. Finally, our models suggest that prefrontal cortex plays a specific and important role in language processing. This has several implications. First, it suggests that other disorders which involve prefrontal cortex (e.g., neurologic patients with lesions of this area) should show language deficits of the sort we have described. It also suggests that prefrontal activity should correlate with performance on language tasks which rely heavily on the processing of context. This represents an exciting area for future research.

Before concluding this section, we should point out that our models have not yet been directly applied to the A \overline{B} task or the Wisconsin Card Sort Test, both of which have been traditionally associated with frontal lobe function and, in the case of the latter, schizophrenic deficits. These tasks (especially the WCST) involve processes of problem solving and hypothesis testing that are not captured by our models in their present form. Nevertheless, they suggest an interpretation of frontal deficits on these tasks that could, in principle, be captured in a simulation model. As we noted earlier, efficient performance in both of these tasks requires that subjects overcome the tendency to repeat response patterns that were correct on previous trials. Thus, both demand that context (e.g., placement of the object on the current trial) be used to control a response tendency (return to prior location) that has gained strength over the course of previous trials. Failure to do so would result in the patterns of perseveration observed (A \overline{B} errors, or failure to switch sorting principle in the WCST). The difference between these tasks and the tasks with response strength asymmetries that we have simulated (Stroop and lexical disambiguation) is that, in the A \overline{B} and WCST, response strength asymmetries develop *within* the task, rather than existing *a priori*. If, however, training of the response pathways was allowed to occur during task performance, then experience on previous trials could lead to the development of response strength asymmetries that could then compete with recent contextual information to determine the response. In fact, Dehaene and Changeux (1989) have proposed a network model of behavioral and electrophysiological data in delayed response tasks, including the A \overline{B} task. This model exhibits principles that are similar, in important respects,

to what we have just described. In their model, a low-level associational module is responsible for mapping stimuli onto responses, and is subject to training on each experimental trial. A higher level module — which can memorize task conditions or perform rule induction — selects or modulates actions performed by the lower level. In this model, the higher level module is assumed to perform the function of the prefrontal cortex. A \bar{B} -type errors arise when this module is impaired, and responses are governed to a greater degree by the training experience of the low-level association module. The similarities between our models and the ones these authors have described — developed independently and with regard to different empirical phenomena — lend strong support to the generality of the principles involved.

5. Biological Disturbances in Schizophrenia

Is dopamine increased or decreased in schizophrenia? We have argued that certain cognitive deficits in schizophrenics can be explained by a *reduction* of dopamine activity in frontal cortex. This may seem to be at odds with what is known about the effects of antipsychotic (neuroleptic) medications. As we discussed above, neuroleptics that tend to improve thought disorder also improve performance on cognitive tasks. For example, performance on the CPT improves with long-term neuroleptic therapy (Spohn et al., 1977), and R. M. Cohen, Semple, Gross, Nordahl et al. (1988) showed that the correlation between prefrontal activity and CPT performance was restored in schizophrenic subjects treated with antipsychotic medications. Yet neuroleptic medications are commonly thought to reduce dopamine activity, by blocking its post-synaptic effects (e.g., Snyder, Banerjee, Yamamura & Greenberg, 1974). This would seem to contradict our hypothesis, which postulates a *reduction* of dopamine tone in frontal cortex. Evidence gathered over the last decade, however, suggests a reconciliation of these points of view. Studies of the effect of neuroleptics on dopamine synthesis have suggested that the mesolimbic and mesocortical dopamine systems respond differently to chronic administration of these medications (for a review see Bannon, Freeman, Chiodo, Bunney & Roth, 1987). These have shown — in rodents, primates and humans — that tolerance to activation of synthesis develops rapidly in the striatal and limbic areas whereas it develops slowly and remains limited in frontal cortex (Scatton, 1977; Scatton, Boireau, Garret, Glowinski & Julou, 1977; Bacopoulos, Spokes, Bird & Roth, 1979; Roth, Bacopoulos, Bustos, and Redmond, 1980). Moreover, during chronic administration of neuroleptics, most dopamine cells enter a state of depolarization inactivation. However, a small number of cells remain active, and the majority of these have been identified as mesocortical cells projecting to frontal cortex (Chiodo & Bunney, 1983). Overall, these data suggest that dopamine tone in prefrontal areas is less affected by neuroleptics than limbic and striatal dopamine. The net result of neuroleptic administration might actually be to enhance dopamine activity in the frontal cortex, at least relative to its activity in other brain regions. This would lead us to expect that neuroleptics would, at worst, have no influence on the cognitive deficits we have addressed and, at best, might lead to improvements.¹¹

¹¹ To the extent that dopamine projections to areas other than frontal cortex affect the pathways which mediate the competing responses in a task, a decrease in dopamine in these pathways will reduce the contrast between these responses. If, at the same time, dopamine in frontal cortex is spared, or enhanced relative to the pathways mediating the responses, then there will be an overall enhancement in the effects of context.

How specific are frontal deficits to schizophrenia? We have argued that a reduction of dopamine in frontal cortex reduces the dynamic range of units in this area. We have not yet explored differences in the behavioral effects of this disturbance and those resulting from other possible disturbances, such as the actual loss of units (corresponding to neurologic lesions of frontal cortex). Indeed, it is possible that there are no differences, and that symptoms of frontal lobe dysfunction in schizophrenia are identical to those of neurologic damage, or other disorders involving dopamine and the frontal lobes (e.g., Parkinson's disease). Nevertheless, to the extent that the models accurately characterize the behavioral consequences of frontal lobe dysfunction, they help delimit the scope of schizophrenic symptoms that can be accounted for on this basis. In so doing, they help to identify findings that can not be related to frontal lobe dysfunction, and for which other explanations must be sought.

Aren't other biological systems involved in schizophrenia? There is little doubt that disturbances of systems other than the frontal lobes are involved in schizophrenia. Other brain regions have been implicated (such as the hippocampus — e.g., Kovelman & Scheibel, 1984; and various subcortical structures — e.g., Crosson & Hughes, 1987, Early, Reiman, Raichle & Spitznagel, 1987, and Stevens, 1973), as have neurotransmitters other than dopamine (such as norepinephrine — e.g., Lake et al., 1980 and van Kammen et al., 1989; and serotonin — e.g., Geyer & Braff, 1987). In their present form, our models are limited in the scope of biological systems that they address. However, we hope that they provide a useful example of how important features of biological processes can be captured within the connectionist framework, and how these can be related to specific behavioral phenomena.

6. Comparison with Other Models of Schizophrenia

A plethora of theories have been proposed to account for the cognitive and biological abnormalities observed in schizophrenia. Here, we focus on those that are most directly related to our own — either by methodology or claims — and that help delineate the specific contributions of our approach.

Broadbent's attentional filter and its breakdown in schizophrenia. Perhaps the most common theory of cognitive dysfunction in schizophrenia draws upon the filter model of attention first proposed by Broadbent (1958; 1971). According to this model, multiple stimuli are registered by the sensory organs and enter a short term store. At this point, stimuli are passed through a filter that provides access to a limited-capacity channel in which further processing takes place. The filter is set by past experience (e.g., conditional probabilities based on past events) and feedback provided by processing in the limited-capacity channel. Investigators who have focused on the phenomenology of schizophrenia (e.g., McGhie & Chapman, 1961; McGhie, 1970; Lang & Buss, 1965; Garnezy, 1977), have suggested that patients experience a difficulty in screening out irrelevant stimuli, and that this may be due to a break-down in the filtering mechanism. Schizophrenics would thus experience one of two states: either a state of stimulus overload in which all stimuli gain equal access to the limited-capacity channel, or a shut-down of information intake in which all stimuli are equally blocked from accessing that channel.

Our models can be related to this conception in several ways. First, the models provide an explicit set of mechanisms for stimulus selection and access to response systems. However, there is no dedicated filter in these models. Rather, a filter-like effect emerges from the interaction of stimulus processing with processing of context when both are channelled through

a common layer of intermediate (associative) units. The models suggest how this filtering of incoming information may be implemented in neural structures. Finally, the models identify the influences that catecholaminergic systems may have on the selection of information, and the consequences of their disruption. Specifically, the models demonstrate that weakening the top-down source of stimulus selection (i.e., the context representation) by reducing dopaminergic tone to the frontal cortex does not result in a complete disorganization of stimulus processing. Rather, degradation follows a distinctive pattern, in which stronger responses begin to dominate weaker ones, regardless of context provided by the task.

Joseph, Frith, and Waddington (1979). These authors describe a mathematical network model which, like our own, relates neural function to higher cognition. They focused on neurotransmitter interactions presumed to support attentional functions. Their model assumes that the dopamine system itself acts as a filter for external inputs and shows how excessive dopamine activity results in exhaustion of inhibitory mechanisms and ultimately a breakdown of filtering functions. This model demonstrates how a neural network can be constructed which performs a filtering function on the basis of simple excitatory and inhibitory interactions. However, Joseph et al. do not relate disturbances of this filtering mechanism to schizophrenic performance in specific behavioral tasks. Because of this, it is difficult to evaluate this model's ability to explain quantitative aspects of cognitive performance.

Hoffman (1987). In this paper, Hoffman reports on a set of computer simulations that display behaviors which are considered to be analogous to several of the positive symptoms of schizophrenia (loosening of associations, blocking and hallucinations). The simulations used fully interconnected Hopfield-type networks as a model of human associative memory processes. During the training phase, the network was taught a set of associations. In the test phase, an input state was specified by activating a subset of the processing units. The network was then allowed to cycle until it settled into a stable configuration of activations. This end-state represented the memory that was accessed from the input specification; this was based on the pattern of connections between units that was learned during training. Hoffman showed that when such a network was forced to encode an excessive number of associations ("memory overload"), specific disturbances of processing occurred: the system often settled into memory states that were inappropriate given the input ("hallucinations") or into states that did not correspond to any of the previously encoded associations ("loosening of associations"). Thus, these simulations related the positive symptoms in schizophrenia to a specific disturbance in the computational mechanisms of the model. Hoffman suggested that this disturbance — memory overload — may arise in schizophrenics as a consequence of a reduced neuronal mass in the prefrontal cortex. In this respect, Hoffman's model can be considered complementary to those we have presented, addressing a different set of symptoms and pathophysiological processes relevant to schizophrenia. As with the model suggested by Joseph et al., however, these models have not yet been applied to the simulation of quantitative aspects of behavior. No doubt, this is due in part to the complex and often inaccessible nature of positive symptoms, which pose serious difficulties for quantification. Indeed, this remains a challenge for all approaches to research on the positive symptoms of schizophrenia.

Weinberger and Berman (1988) and Levin (1984). Both of these groups have marshalled empirical support for the involvement of frontal lobe dysfunction in schizophrenia. Furthermore, these investigators have specifically suggested that a deficit in the dopaminergic innervation of the frontal cortex is responsible for performance impairments in tasks such as the WCST, and, from a more clinical perspective, for the negative symptoms of schizophrenia. Our models agree with this point of view; they also extend it in several important ways. First, they go beyond earlier hypotheses by proposing a specific set of mechanisms which explain the relationship between a disturbance in dopamine activity, frontal lobe function, and task

performance. This has allowed us to address quantitative aspects of performance in a number of behavioral tasks, and to provide a unified account of schizophrenic patterns of performance in terms of a common underlying deficit. Although we have not yet applied our models to performance on the Wisconsin Card Sort Test, we discussed how schizophrenic deficits on this task could be related to a disturbance in the processing of context, and suggested how the models could be extended to test this idea.

7. The Role of Computational Modelling

The question is often asked: How do models contribute to an understanding of the data they simulate? After all, the data already exist, and the principles or ideas captured by a model can often be expressed without the use of a computer program (indeed, some would contend that this must be so if the ideas are of any general value). McClelland (1988) has provided an articulate reply to this question, in describing the relevance of models to empirical investigations in psychology. He points out that models can: a) bring seemingly disparate empirical phenomena together under a single explanation; b) provide new interpretations of existing findings; c) reconcile contradictory evidence; and d) lead to new predictions. Throughout the present discussion, we have tried to show how our models realize these different goals. For example, the models identified a disturbance in the processing of context that could explain impairments of attention, language processing and overall reaction time in schizophrenia; they revealed that an overall increase in reaction time could arise from a specific rather than a generalized information processing deficit; they suggested a reconciliation of contradictory findings with respect to the CPT and prefrontal activation; and they led to predictions concerning normal and schizophrenic performance on behavioral tasks, as well as predictions about dopamine effects on prefrontal metabolism. McClelland also emphasizes the role that models play in formalizing theoretical concepts. By committing a set of ideas to a computer program, and examining their ability to account for quantitative data, the ideas are put to a rigorous test of both their internal coherence and the resolution of their explanatory power.

Most important, however, is the role that modelling plays in the discovery process. At times the insights provided by a model may seem, in hindsight, to be obvious or not to have required the effort involved in constructing a computer simulation. Usually, however, such a perception fails to recognize that the insight came from the process of developing the model itself. The three models described in this paper were actually developed independently, and for different purposes. The Stroop model was developed to account for normal performance in this task; the CPT simulation was developed to explore gain as a model of catecholaminergic effects on behavior; and the lexical disambiguation model was developed specifically to address schizophrenic language deficits. It was only when we compared the mechanisms at work in these different models that we realized all relied on common principles of processing. This, in conjunction with our work with the gain parameter in the CPT model, suggested a hypothesis about the relationship between biological and behavioral factors in schizophrenia. In this way, the models provided an important vehicle for the discovery — and not just the testing — of new ideas.

V. Conclusion

We have tried to show how the connectionist framework can be brought to bear on the relationship between some of the biological and cognitive disturbances characteristic of schizophrenia. The models we have presented suggest that a common information processing deficit underlies impaired performance in attention and language processing tasks. The models related this deficit to decreased dopaminergic activity in prefrontal cortex. The models, and the simulations based on them relied on many simplifying assumptions and provided, at best, a coarse approximation of the mechanisms underlying both normal and schizophrenic behavior. While accounting for empirical data is a primary goal in the development of computer simulation models, McClelland (1988) has argued that this may not be the only basis for their evaluation. Models are useful if they offer new interpretations of empirical phenomena, unify previously unrelated observations, reconcile conflicting findings, and predict new empirical facts. We have indicated how our models — simple as they are — may fulfill these different functions. In so doing, we hope that these models will help provide a more refined and integrated approach to the riddle of behavioral and biological disturbances in schizophrenia.

Authors' Notes

The order of authorship is alphabetical; both authors made equal contributions to the research and preparation of this manuscript. A preliminary version of this work was presented at the 142nd Annual meeting of the American Psychiatric Association, May, 1989.

We gratefully acknowledge the helpful comments and suggestions made by David Galin, Steve Matthyse, James McClelland and Benoit Mulsant in the development of our ideas and the preparation of this paper. In addition John Csernansky, Len Horowitz, Roy King, Tracy Kristoffersen, David Spiegel, Elisabeth Targ, and Sue Theimann participated in the design and implementation of the experimental study of schizophrenic language performance that we reported.

This work was supported by a NIMH Physician Scientist Award (MH00673) to the first author and a NIMH Individual Fellow Award (MH09696) to the second author. Part of this work was also supported by a research grant from the Scottish Rite Schizophrenia Research Program, N.M.J., U.S.A. to the first author, and a research fellowship from Western Psychiatric Institute and Clinic to the second author.

Correspondence concerning this paper should be addressed to Jonathan D. Cohen (Department of Psychology) or David Servan-Schreiber (School of Computer Science), Carnegie Mellon University, Pittsburgh, PA 15213.

References

- Abramczyk RR, Jordan DE & Hegel M (1983). "Reverse" Stroop effect in the performance of schizophrenics *Perceptual and Motor Skills* 56:99-106.
- Andreasen N, Nasrallah HA, Dunn V, Olson SC, Grove WM, Ehrhardt JC, Coffman JA & Crossett JHW (1986). Structural abnormalities in the frontal system of schizophrenia. *Archives of General Psychiatry* 43:136-144.
- Aou S, Nishino H, Inokuchi A & Mizuno Y (1983) Influence of catecholamines on reward-related neuronal activity in monkey orbitofrontal cortex. *Brain Research* 267:165-170
- Asarnow RF & MacCrimmon (1978). Residual performance deficit in clinically remitted schizophrenics: A marker of schizophrenia? *Journal of Abnormal Psychology* 87:597-608.
- Bacopoulos NG, Spokes EG, Bird ED & Roth RH (1979). Antipsychotic drug action in schizophrenic patients: Effects on cortical dopamine metabolism after long-term treatment. *Science* 205:1405-1407.
- Bannon MH, Freeman AS, Chiodo LA, Bunney BS & Roth RH (1987). The electrophysiological and biochemical pharmacology of the mesolimbic and mesocortical dopamine neurons. In LI Iversen, SD Iversen, and SH Snyder (Eds.), *Handbook of Psychopharmacology Vol 19* (pp. 329-374). New York: Plenum Press.
- Benjamin TB & Watt NF (1969). Psychopathology and semantic interpretation of ambiguous words. *Journal of Abnormal Psychology* 74(6):706-714.
- Berman KF, Illowsky BP & Weinberger DR (1988). Physiological dysfunction of dorsolateral prefrontal cortex in schizophrenia: Further evidence for regional and behavioral specificity. *Archives of General Psychiatry* 45:616-622.
- Berman KF, Zec RF & Weinberger DR (1986). Physiological dysfunction of dorsolateral prefrontal cortex in schizophrenia: II. Role of neuroleptic treatment, attention and mental effort. *Archives of General Psychiatry* 43:126-135.
- Bird ED, Barnes J, Iversen LL, Spokes EG, Mackay AVP & Shepherd M (1977). Increased brain dopamine and reduced glutamine acid decarboxylase and choline acetyl transferase activity in schizophrenia and related psychoses. *Lancet* 2:1157-1158.
- Blanley PH (1974). Two studies on the language behavior of schizophrenics. *Journal of Abnormal Psychology* 83(1):23-31.
- Bloom FE, Schulman JA & Koob GF (1989). Catecholamines and Behavior. In U Trendelenburg and N Wiener (Eds.), *Handbook of Experimental Pharmacology, (Ch. 11)* 90:27-88 Berlin: Springer Verlaag.
- Borst U & Cohen R (1989). Filling the preparatory interval: crossover effect. *Psychiatric Medicine*, in press.
- Bowen FP, Kammienny RS, Burns MM & Yahr MD (1975). Parkinsonism: Effects of L-dopa treatment on concept formation. *Neurology* 25:701-704.
- Bowers MB Jr, Heninger GR & Sternberg D (1980). Clinical processes and central dopaminergic activity in psychotic disorders. *Communications in Psychopharmacology* 4:177-188.

- Broadbent DE (1958). *Perception and communication*. Oxford: Pergamon Press.
- Broadbent DE (1971). *Decision and Stress*. New York: Academic Press.
- Brozoski TJ, Brown RM, Rosvold HE, & Goldman PS (1979). Cognitive deficit caused by regional depletion of dopamine in prefrontal cortex of Rhesus monkey. *Science* 205:929-931.
- Buchsbaum MS, Ingvar DH, Kessler R, Waters RN, Cappelletti J, van Kammen DP, King CA, Johnson JL, Manning RG, Flynn RW, Mann LS, Bunney WE & Sokoloff L (1982). Cerebral glucography with positron tomography. *Archives of General Psychiatry* 39:251-259.
- Burnod Y & Korn H (1989) Consequences of stochastic release of neurotransmitters for network computation in the central nervous system. *Proceedings of the National Academy of Science (U.S.A)* 86:352-356
- Chapman LJ & Chapman JP (1978). The measurement of differential deficit. *Journal of Psychiatric Research* 14:303-311.
- Chapman LJ, Chapman JP & Miller GA (1964). A theory of verbal behavior in schizophrenia. In Maher BA (Ed.), *Progress in Experimental Personality Research, Vol 1*. New York: Academic Press.
- Chiodo LA & Berger TW (1986). Interactions between dopamine and amino acid-induced excitation and inhibition in the striatum. *Brain Research* 375:198-203.
- Chiodo LA & Bunney BS (1983). Typical and atypical neuroleptics: Differential effects of chronic administration on the activity of A9 and A10 midbrain dopaminergic neurons. *Journal of Neuroscience* 3:1607-1619.
- Cohen BM (1981). Dopamine receptors and antipsychotic drugs. *McLean Hospital Journal* 1(2):95-115.
- Cohen JD, Dunbar K & McClelland JL (1990). On the control of automatic processes: A parallel distributed processing model of the Stroop effect. *Psychological Review*, in press.
- Cohen JD, Targ E, Kristoffersen T & Spiegel D (1989). The fabric of thought disorder: Disturbances in the processing of context. In submission; available upon request.
- Cohen R, Borst U & Rist F (1989). Cross-over and cross-modality effects in schizophrenia: Some old and some new data. In submission.
- Cohen RM, Semple WE, Gross M, Holcomb HH, Dowling MS & Nordahl TE (1988). Functional localization of sustained attention: Comparison to sensory stimulation in the absence of instruction. *Neuropsychiatry, Neuropsychology, and Behavioral Neurology* 1:2-20.
- Cohen RM, Semple WE, Gross M, Nordahl TE, DeLisi LE, Holcomb HH, King AC, Morihisa JM & Pickar D (1987). Dysfunction in a prefrontal substrate of sustained attention in schizophrenia. *Life Sciences* 40:2031-2039.
- Cohen RM, Semple WE, Gross M, Nordahl TE, Holcomb HH, Dowling MS & Pickar D (1988). The effect of neuroleptics on dysfunction in a prefrontal substrate of sustained attention in schizophrenia. *Life Sciences* 43:1141-1150.
- Cozolino LJ (1983). The oral and written productions of schizophrenic patients. In Maher BA (Ed.), *Progress in Experimental Personality Research, Vol. 12*. New York: Academic Press.

- Creese I, Burt DR & Snyder SH (1976). Dopamine receptor binding predicts clinical and pharmacological potencies of antischizophrenic drugs. *Science* 192:481-483.
- Cross AJ, Crow TJ & Owen F (1981). 3-H-flupenthixol binding in post-mortem brains of schizophrenics: Evidence for a selective increase in dopamine D2 receptors. *Psychopharmacology* 74:122-124.
- Crosson B & Hughes CW (1987). Role of the thalamus in language: is it related to schizophrenic thought disorder? *Schizophrenia Bulletin* 13(4):605-621.
- Crow TJ (1980). Molecular pathology of schizophrenia: More than one disease process? *British Medical Journal* 137:383-386.
- Crow TJ, Cross AJ, Johnson JA, Johnstone EC, Joseph MH, Owen F, Owens DGC & Poulter M (1984). Catecholamines and schizophrenia: An assessment of the evidence. In A. R. Liss (Ed.), *Catecholamines: Neuropharmacology and central nervous system-therapeutic aspects*.
- Dehaene S & Changeux JP (1989). A simple model of prefrontal cortex function in delayed-response tasks. *Journal of Cognitive Neuroscience* 1:244-261.
- Diamond A & Doar B (1989). The performance of human infants on a measure of frontal cortex function, the delayed response task. *Developmental Psychobiology* 22(3):271-294.
- Diamond A (1989a). Developmental progression in human infants and infant monkeys, and the neural bases of, A B and delayed response performance. In Diamond A (Ed.), *The development and neural bases of higher cognitive functions*. New York: NY Academy of Science Press.
- Diamond A (1989b). Infants' and young children's performance on delayed non-match to sample (direct and indirect) and visual paired comparison. In Diamond A (Ed.), *The development and neural bases of higher cognitive functions*. New York: NY Academy of Science Press.
- Diamond A (1989c). Developmental progression in human infants and infant monkeys, and the neural bases of, inhibitory control of reaching. In Diamond A (Ed.), *The development and neural bases of higher cognitive functions*. New York: NY Academy of Science Press.
- Diamond, A (1985). Developmental of the ability to use recall to guide action, as indicated by infants' performance on AB. *Child Development* 56:868-883.
- Dunbar, K., & MacLeod, C.M. (1984). A horse race of a different color: Stroop interference patterns with transformed words. *Journal of Experimental Psychology: Human Perception and Performance* 10:622-639.
- Dyer, F.N. (1973). The Stroop phenomenon and its use in the study of perceptual, cognitive, and response processes. *Memory and Cognition* 1:106-120.
- Early TS, Reiman EM, Raichle ME & Spitznagel EL (1987). Left globus pallidus abnormality in never-medicated patients with schizophrenia. *Proceedings of the National Academy of Science U.S.A.* 84:561-563.
- Erlenmeyer-Kimling L & Cornblatt B (1978). Attentional measures in a study of children at high-risk for schizophrenia. *Journal of Psychiatric Research* 14:93-98.
- Farkas T, Wolf AP, Jaeger J, Brodie JD, Christman DR & Fowler JS (1984). Regional brain glucose metabolism in chronic schizophrenia. *Archives of General Psychiatry* 41:293-300.

- Ferron A, Thierry AM, Le Douarin C & Glowinski J (1984) Inhibitory influence of the mesocortical dopaminergic system on spontaneous activity or excitatory response induced from the thalamic mediodorsal nucleus in the rat medial prefrontal cortex. *Brain Research* 302:257-265.
- Foote SL, Freedman R & Oliver AP (1975). Effects of putative neurotransmitters on neuronal activity in monkey auditory cortex. *Brain Research* 86:229-242.
- Franzen G & Ingvar DH (1975). Abnormal distribution of cerebral activity chronic schizophrenia. *Journal of Psychiatric Research* 12:199-214.
- Freeman WJ (1979). Nonlinear gain mediating cortical stimulus-response relations. *Biological Cybernetics* 33:243-247.
- Fuster JM (1980). *The prefrontal cortex*. Raven Press.
- Garnezy N (1977). The psychology and psychopathology of attention. *Schizophrenia Bulletin* 3:360-369.
- Geraud G, Arne-Bes MC, Guell A & Bes A (1987). Reversibility of hemodynamic hypofrontality in schizophrenia. *Journal of Cerebral Blood Flow and Metabolism* 7:9-12.
- Geyer MA & Braff DL (1987). Startle habituation and sensorimotor gating in schizophrenia and related animal models. *Schizophrenia Bulletin* 13(4):643-688.
- Goldman-Rakic PS (1985). Circuitry of primate prefrontal cortex and regulation of behavior by representational memory. *Handbook of Physiology - The Nervous System* 5:373-417.
- Grant DA & Berger EA (1948). A behavioral analysis of degree of reinforcement and ease of shifting to new responses in a Weigl type card sorting problem. *Journal of Experimental Psychology* 38:404-411.
- Green DM & Swets JA (1966). *Signal detection theory and psychophysics*. Huntington, NY: Robert E. Krieger Publishing Co.
- Gur RE, Gur RC, Skolnick BE, Caroff S, Obrist WD, Resnick S & Reivich M (1985). Brain function in psychiatric disorders: III. Regional cerebral blood flow in unmedicated schizophrenics. *Archives of General Psychiatry* 42:329-334.
- Gur RE, Skolnick BE, Gur RC, Caroff S, Riegar W, Obrist WD, Younkin D & Reivich M (1983). Brain function in psychiatric disorders: I. Regional cerebral blood flow in medicated schizophrenics. *Archives of General Psychiatry* 40:1250-1254.
- Harvey PD (1983). Speech competence in manic and schizophrenic psychoses: The association between clinically rated thought disorder and cohesion and reference performance. *Journal of Abnormal Psychology* 92(3):368-377.
- Hinton GE & Shallice T (1989). Lesioning a connectionist network: Investigations of acquired dyslexia. Technical Report CRG-TR-89-3, Department of Computer Science, University of Toronto, 10 Kings College Road, Toronto M5S 1A4, Canada.
- Hoffman RE (1987). Computer simulations of neural information processing and the schizophrenia-mania dichotomy. *Archives of General Psychiatry* 44:178-188.
- Hu XT & Wang RY (1988) Comparison of effects of D₁ and D₂ dopamine receptor agonists on neurons in the rat caudate putamen: An electrophysiological study. *The Journal of Neuroscience* 8:4340-4348.
- Ingvar DH & Franzen G (1974). Abnormalities of cerebral flow distribution in patients with chronic schizophrenia. *Acta Psychiatr Scand* 50:425-462.

- Janowsky DS, & Rich C (1979). Amphetamine psychosis and psychotic symptoms. *Psychopharmacology (Berlin)* 65:73-77.
- Janowsky DS, Huey L, Storms L & Judd LL (1977). Methylphenidate hydrochloride effects on psychological tests in acute schizophrenic and nonpsychotic patients. *Archives of General Psychiatry* 34:189-194.
- Joseph MH, Frith CD & Waddington JL (1979). Dopaminergic mechanisms and cognitive deficits in schizophrenia. *Psychopharmacology (Berlin)* 63:273-280.
- Kolb B & Whishaw IQ (1983). Performance of schizophrenic patients on tests sensitive to left or right frontal, temporal, or parietal function in neurologic patients. *Journal of Mental and Nervous Disease* 171:435-443.
- Kopfstein JH & Neal JM (1972). A multivariate study of attention dysfunction in schizophrenia. *Journal of Abnormal Psychology* 3:294-298.
- Kornetsky C & Orzack MH (1964). A research note on some of the critical factors on the dissimilar effects of chlorpromazine and secobarbital on the digit symbol substitution and continuous performance tests. *Psychopharmacologia* 6:79.
- Kornetsky C (1972). The use of a simple test of attention as a measure of drug effects in schizophrenic patients. *Psychopharmacologia (Berl.)* 24:99-106.
- Kovelman JA & Scheibel AB (1984). A neurohistological correlate of schizophrenia. *Biological Psychiatry* 19:1601-1621.
- Kraepelin E (1950). In Zinkin J (trans.), *Dementia Praecox and paraphrenia*. New York: International Universities Press Inc.
- Lake CR, Sternberg DE, van Kammen DP, Ballenger JC, Ziegler MG, Post RM, Kopin IJ & Bunney WE (1980). Schizophrenia: Elevated cerebrospinal fluid norepinephrine. *Science* 207:331-333.
- Lang PJ & Buss AH (1965). Psychological deficit in schizophrenia : II. Interference and activation. *Journal of Abnormal Psychology* 70:77-106.
- Lehky SR & Sejnowski TJ (1988). Network model of shape-from-shading: neural function arises from both receptive and projective fields. *Nature* 333:452-454.
- Levin S (1984). Frontal lobe dysfunctions in schizophrenia — II. Impairments of psychological and brain functions. *Journal of Psychiatry Research* 18(1):57-72.
- Levin S, Yurgelun-Todd D & Craft S (1989). Contributions of clinical neuropsychology to the study of schizophrenia. *Journal of Abnormal Psychology*, in press.
- Luria AR (1966). *Higher cortical functions in man*. London: Tavistock.
- MacLeod CM (1989). Half a century of research on the Stroop effect: An integrative review. Manuscript submitted for publication.
- Maher BA (1972). The language of schizophrenia: A review and interpretation. *British Journal of Psychiatry* 120:3-17.
- Malmo HP (1974). On frontal lobe functions: Psychiatric patient controls. *Cortex* 10:231-237.
- McClelland JL & Rumelhart DE (Eds.). (1986). *Parallel Distributed Processing Explorations in the Microstructure of Cognition*. Vol. 2: *Psychological and biological models*. Cambridge, MA: MIT Press.

- McClelland JL (1988). Connectionist models and psychological evidence. *Journal of Memory and Language* 27:107-123.
- McCulloch J, Savaki HE, McCulloch MC, Jehle J & Sokoloff L (1982). The distribution of alterations in energy metabolism in the rat brain produced by apomorphine. *Brain Research* 243:67-80.
- McGhie A (1970). Attention and perception in schizophrenia. In BA Maher (Ed.), *Progress in Experimental Personality Research*, Vol 5 (pp. 1-36). New York: Academic Press.
- McGhie A & Chapman J (1961). Disorders of attention and perception in early schizophrenia. *Brit. J. of Med. Psychol.* 34:103-116.
- Meltzer HY & Stahl SM (1976). The dopamine hypothesis of schizophrenia: a review. *Schizophrenia Bulletin* 2:19-76.
- Milner B (1963). Effects of different brain lesions on card sorting. *Archives of Neurology* 9:90-100.
- Morihisa JM & McAnulty GB (1985). Structure and function: Brain electrical activity mapping and computed tomography in schizophrenia. *Biological Psychiatry* 20:3-19.
- Nauta JHW & Domesick VB (1981). Ramifications of the limbic system. In Matthysse S, (Ed.), *Psychiatry and the Biology of the Human Brain: A Symposium Dedicated to Seymour S. Kety* (pp.165-188). New York: Elsevier/North-Holland.
- Neale JM (1971). Perceptual span in schizophrenia. *Journal of Abnormal Psychology* 77:196-204.
- Nelson HE (1976). A modified card sorting test sensitive to frontal lobe defects. *Cortex* 12:313-324.
- Nuechterlein KH (1977). Reaction time and attention in schizophrenia: A critical evaluation of the data and theories. *Schizophrenia Bulletin* 3:373-428.
- Nuechterlein KH (1984). Information processing and attentional functioning in the developmental course of schizophrenic disorders. *Schizophrenia Bulletin* 10:160-203.
- Nuechterlein KH (1983). Signal detection in vigilance tasks and behavioral attributes among offspring of schizophrenic mothers and among hyperactive children. *Journal of Abnormal Psychology* 92:4-28.
- Oades RD (1981). Dopaminergic agonist and antagonistic drugs in the ventral tegmentum or rats inhibit and facilitate changes of food search behaviour. *Neurosci Lett* 27:75-80.
- Parasuraman R (1979). Memory load and event rate control sensitivity decrements in sustained attention. *Science* 205:924-927.
- Perret E (1974). The left frontal lobe of man and the suppression of habitual responses in verbal categorical behaviour. *Neuropsychologia* 12:323-330.
- Piaget J (1954). *The origins of intelligence in children*. New York: Basic Books. Original French edition, 1937.
- Pickar D, Labarca R, Linnoila M, Roy A, Hommer D, Everett D & Paul SM (1984). Neuroleptic-induced decrease in plasma homovanillic acid and antipsychotic activity in schizophrenic patients. *Science* 225:954-957.
- Reader TA, Ferron A, Descarries L & Jasper HH (1979) Modulatory role for biogenic amines in the cerebral cortex: microiontophoretic studies. *Brain Research* 160:217-229

- Robinson AL, Heaton RK, Lehman RAW & Stilson DW (1980). The utility of the Wisconsin Card Sorting Test in detecting and localizing frontal lobe lesions. *Journal of Consulting Clinical Psychology* 48:605-614.
- Rochester SR & Martin JR (1979). *Crazy talk: A study in the discourse of schizophrenic speakers*. New York: Plenum.
- Rolls ET, Thorpe SG, Boytim M, Szabo I & Perrett DI (1984). Responses of striatal neurons in the behaving monkey. 3. Effects of iontophoretically applied dopamine on normal responsiveness. *Neuroscience* 12:1201-1212.
- Rosvold HE, Mirsky AF, Sarason I, Bransome ED & Beck LH (1956). A continuous performance test of brain damage. *Journal of Consulting Psychology* 20(5):343-350.
- Roth RH, Bacopoulos NG, Bustos G, and Redmond DE (1980). Antipsychotic drugs: Differential effects on dopamine neurons in basal ganglia and mesocortex following chronic administration in human and non-human primates. In F Cattabeni, G Racagni, PF Spano & E Costa (Eds.), *Advances in Biochemical Psychopharmacology Vol. 24* (pp. 513-520). New York: Raven Press.
- Rumelhart DE & McClelland JL (1986b). PDP models and general issues in cognitive science. In DE Rumelhart & JL McClelland (Eds.), *Parallel Distributed Processing: Explorations in the microstructure of cognition. Vol. 1: Foundations* (pp. 110-146). Cambridge, MA: MIT Press.
- Rumelhart DE & McClelland JL (1986c). On learning the past tense of English verbs. In JL McClelland & DE Rumelhart (Eds.), *Parallel Distributed Processing: Explorations in the microstructure of cognition. Vol 2: Psychological and biological models* (pp. 216-271). Cambridge, MA: MIT Press.
- Rumelhart DE, Hinton GE & Williams RJ (1986). Learning internal representations by backpropagating errors. *Nature* 323:533-536.
- Rumelhart DE, McClelland JL (Eds.) (1986a). *Parallel Distributed Processing Explorations in the Microstructure of Cognition. Vol. 1: Foundations*. Cambridge, MA: MIT Press.
- Rutschmann J, Cornblatt B & Erlenmeyer-Kimling L (1977). Sustained attention in children at risk for schizophrenia. *Archives of General Psychiatry* 34:571-575.
- Rutter DR (1979). The reconstruction of schizophrenic speech. *British Journal of Psychiatry* 134:356-359.
- Salzinger K (1971). An hypothesis about schizophrenic behavior. *American Journal of Psychotherapy* 25:601-614.
- Salzinger K, Portnoy S & Feldman RS (1964). Verbal behavior of schizophrenic and normal subjects. *Annals of the New York Academy of Sciences*. 105:845-860.
- Salzinger K, Portnoy S, Pisoni DB & Feldman RS (1970). The immediacy hypothesis and response-produced stimuli in schizophrenic speech. *Journal of Abnormal Psychology* 76(2):258-264.
- Sawaguchi T & Matsumura M (1985) Laminar distributions of neurons sensitive to acetylcholine, noradrenaline and dopamine in the dorsolateral prefrontal cortex of the monkey. *Neuroscience Research* 2:255-273
- Scatton B (1977). Differential regional development of tolerance to increase in dopamine turnover upon repeated neuroleptic administration. *European Journal of Pharmacology* 46:363-369.

- Scatton B, Boireau A, Garret C, Glowinski J & Julou L (1977). Action of the palmitic ester of pipotiazine on dopamine metabolism in the nigro-striatal, meso-limbic and meso-cortical systems. *Naunyn-Schmiedeberg's Archives of Pharmacology* 296:169-175.
- Schwartz S (1982). Is there a schizophrenic language? *The Behavioral and Brain Sciences* 5:579-626.
- Sedvall G, Fyro B, Nyback H, Wiesel FA and Wode-Helgodt B (1974). Mass fragmentometric determination of HVA in lumbar cerebrospinal fluid of schizophrenic patients during treatment with antipsychotic drugs. *Journal of Psychiatric Research* 11:75-80.
- Seidenberg MS & McClelland JL (1990). A distributed development model of word recognition and naming. *Psychological Review*, in press.
- Servan-Schreiber D (1989). Catecholamines and behavior: a connectionist approach to neuromodulation. Unpublished manuscript.
- Shakow D (1962). Segmental set: A theory of the formal psychological deficit in schizophrenia. *Archives of General Psychiatry* 6:1-17.
- Shallice T (1982). Specific impairments of planning. *Phil. Trans. R. Soc. Lond. B* 298:199-209.
- Simon H, Scatton B & Le Moal M (1980). Dopaminergic A10 neurones are involved in cognitive functions. *Nature* 286:150-151.
- Snyder SH (1972). Catecholamines in the brain as mediators of amphetamine psychosis. *Archives of General Psychiatry* 27:169-179.
- Snyder SH (1976). The dopamine hypothesis of schizophrenia: Focus on the dopamine receptor. *American Journal of Psychiatry* 133:197-202.
- Snyder SH, Banerjee SP, Yamamura HI & Greenberg D (1974). Drugs, neurotransmitters and schizophrenia. *Science* 184:1243-1253.
- Spohn HE, Lacoursiere RB, Thomson K & Coyne L (1977). Phenothiazine effects on psychological and psychophysiological dysfunction in chronic schizophrenics. *Archives of General Psychiatry* 34:633-644.
- Spring B (in press). Distractibility as a marker of vulnerability to schizophrenia. *Psychopharmacology Bulletin*.
- Stevens JR (1973). An anatomy of schizophrenia? *Archives of General Psychiatry* 29:177-189.
- Strauss ME (1975). Strong meaning-response bias in schizophrenia. *Journal of Abnormal Psychology* 84(3):293-298.
- Stroop JR (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology* 18:643-662.
- Stuss DT & Benson DF (1984). Neuropsychological studies of the frontal lobes. *Psychological Bulletin* 95:3-28.
- Swets JA & Sewall ST (1963). Invariance of signal detectability over stages of practice and levels of motivation. *Journal of Experimental Psychology* 66:120-126.
- van Kammen DP, Peters J, van Kammen WB, Nugent A, Goetz KL, Yao J & Linnoila M (1989). CSF norepinephrine in schizophrenia is elevated prior to relapse after haloperidol withdrawal. *Biological Psychiatry* 26:176-188.

- Wapner S & Krus DM (1960). Effects of lysergic acid diethylamide, and differences between normals and schizophrenics on the Stroop color-word test. *Journal of Neuropsychiatry* 2:76-81.
- Weinberger DR & Berman KF (1988). Speculation on the meaning of metabolic hypofrontality in schizophrenia. *Schizophrenia Bulletin* 14:157-168.
- Weinberger DR, Berman KF & Chase TN (1988). Mesocortical dopaminergic function and human cognition. *Annals of the New York Academy of Science* (Vol 537: "The mesocorticolimbic dopamine system") (pp. 330-338).
- Weinberger DR, Berman KF & Illowsky BP (1988). Physiological dysfunction of dorsolateral prefrontal cortex in schizophrenia: A new cohort of evidence for a monoaminergic mechanism. *Archives of General Psychiatry* 45:606-615.
- Weinberger DR, Berman KF & Zec RF (1986). Physiological dysfunction of dorsolateral prefrontal cortex in schizophrenia: I. Regional cerebral blood flow evidence. *Archives of General Psychiatry* 43:114-125.
- Weinberger DR, Bigelow LB, Kleinman JE, Klein ST, Rosenblatt JE & Wyatt RJ (1980). Cerebral ventricular enlargement in chronic schizophrenia. *Archives of General Psychiatry* 37:11-13.
- Wielgus MS & Harvey PD (1988). Dichotic listening and recall in schizophrenia and mania. *Schizophrenia Bulletin* 14(4):689-700.
- Zipser D & Andersen RA (1988). A back-propagation programmed network that simulates response properties of a subset of posterior parietal neurons. *Nature* 331:679-684.